

The Margin of Inefficiency

Reading Prediction Market Prices as Bounds on Probability

Harry Crane*

October 2026

Abstract

Prediction market prices are often interpreted as probabilities. Here we discuss when this interpretation is appropriate and when it is not. The key conceptual finding is that a price does not correspond directly to probability, but instead places constraints on the underlying probability. These constraints imply an *inefficiency interval* which contains the true underlying probability under mild assumptions on trader behavior. The width of the inefficiency interval is called the *margin of inefficiency* (MOI). The inefficiency interval is the set of probabilities at which no trade at the quoted prices has positive expected value after the bid-ask spread, fees, and commissions are paid. The margin of inefficiency quantifies the uncertainty of the market price, akin to the margin of error in traditional polling. This paper states the core idea, defines the main concepts, and shows how to combine the margins of several markets on the same event to arrive at a consensus.

1 Introduction

On the morning of the 2020 U.S. presidential election, the contract for Joe Biden to win California traded at \$0.94 on PredictIt. California had not voted for a Republican presidential candidate since 1988, and Biden was expected to win California easily. He went on to win the state by 29 points. Other exchanges priced the same contract at about \$0.99. If the PredictIt price were read as a probability, it would imply that Donald Trump had a 6% chance of winning California, which is too high by any estimate.

In the fall of 2024, PredictIt and Polymarket both listed the presidential election between Donald Trump and Kamala Harris. In the month before the election, PredictIt's closing price for Trump was below Polymarket's on all but two days, by an average of 3.4 points 3.4. On October 28 the price was \$0.59 on PredictIt and \$0.66 on Polymarket. The gap was widely reported, and cited as an incoherence and inefficiency in prediction markets. Some claimed that Polymarket's price was caused by a small number very large pro-Trump traders. Others claimed that PredictIt's price reflected the Democratic-leaning politics of its users.

*Department of Statistics, Rutgers University. Email: harry.crane@rutgers.edu. Member, CFTC Technology Advisory Committee.

These examples raise natural questions about the accuracy and precision of prediction markets for forecasting events. If price correspond to probabilities, then one would expect a market on a near-certain event (e.g., Biden to win the electoral college in California) to trade close to \$1.00, not at \$0.94 on election morning as it did on PredictIt's 2020 market. And, if markets are rational and are efficient at incentivizing participants to trade based on their information, one would not expect two markets on the same event to trade 7 points apart, as PredictIt and Polymarket did in the 2024 presidential election.

Here we reconcile these discrepancies by analyzing how market design introduces observable frictions that prevent price-sensitive traders from correcting obvious deviations from the true underlying probability. In the examples above, PredictIt charged a commission of 10% on the profit of winning contracts, a fee of 5% on withdrawals from their platform, and also limited each trader to \$850 per contract. When considering all of these factors together, even a trader who was 100% certain that Biden would win California would likely not have the incentive to purchase at \$0.94. Given the fees and commissions charged, the contract settles at \$1.00 for a profit of \$0.06, of which \$0.006 would go to commission, and the withdrawal fee of 5% applies to the entire amount (principal plus profit). So a trader who risked the limit of \$850 on this market would have purchased 904 shares, which would be worth \$904 before any commissions or fees are applied. After commission of \$5.40 on the \$54 profit, there would be \$898.60 left, which after the 5% withdraw fee leaves \$853.67, for a profit of \$3.67. And clearly if the price moved higher to \$0.95, the above proposition would be negative. As this demonstrates, the \$0.94 price on PredictIt was a reflection of poor market design and trader rationality, rather than the initial inference that markets are irrational.

Moreover, the market price of \$0.94 is not a claim that Biden's probability was 94%. It said that his probability was no higher than about 99.6%, which is the probability at which buying at \$0.94 stops losing money after charges. The same accounting on the other side of the market gives a lower limit. The price placed the probability in a range, and the range included 99%. The 2024 gap has the same explanation. PredictIt's range for Trump on October 28 was about seventeen points wide, and Polymarket's price was inside it.

The general point is as follows. Price equals probability in the absence of friction. In practice there is always friction, even in the most liquid markets. The accurate statement is that price imposes bounds on probability, and the frictions of the market determine how far apart the bounds are. We call the range between the bounds the *inefficiency interval* and its width the *margin of inefficiency*. A market with a narrow margin locates the probability closely. A market with a wide margin does not, and its price should not be compared with other prices as if it did.

Section 2 states the Fundamental Theorem of Prediction Markets and derives from it the bounds that a price places on probability in a market with no fees. Section 3 defines the inefficiency interval and the margin of inefficiency for a market with fees and commissions, and returns to the California example. Section 4 discusses the frictions that the formula leaves out. Section 5 combines the intervals of several markets on the same event into a consensus, with the 2024 election as the example. Section 6 applies the method to the 2026 midterm elections. Section 7 discusses interpretation.

2 The Fundamental Theorem of Prediction Markets

An *event contract* (often shortened to *contract*) on an outcome pays \$1.00 if the outcome occurs and \$0.00 otherwise. A market on a Yes-or-No event quotes two such contracts, Yes and No. A market is *efficient* if no contract quoted in it has positive expected value. In practical terms, an efficient market cannot be profitably traded. The interpretation of price as probability follows from this definition through the Fundamental Theorem of Prediction Markets [Crane, forthcoming].

Fundamental Theorem of Prediction Markets. A prediction market is efficient if and only if the price of every contract quoted in the market is at least the expected settlement value of that contract. In the absence of friction, the price of a contract on an outcome equals the probability of that outcome.

The theorem is a direct consequence of the expected value calculation. Suppose the outcome occurs with probability p and the contract on it is priced at π . A trader who buys the contract earns $1 - \pi$ with probability p and loses π with probability $1 - p$, for an expected value of

$$p \times (1 - \pi) + (1 - p) \times (-\pi) = p - \pi.$$

Efficiency requires this to be at most zero, so $\pi \geq p$. Now apply the same condition to the opposite contract, which pays when the outcome does not occur. In a market with no friction, the opposite contract costs $1 - \pi$ and its expected value is $(1 - p) - (1 - \pi) = \pi - p$. Efficiency requires $\pi - p \leq 0$, so $\pi \leq p$. Together the two conditions force $\pi = p$. Price equals probability.

The derivation makes the assumptions explicit. The first is that traders act on positive expected value, so that any contract with positive expected value gets bought until its price rises and the expected value is gone. The second is that there is no friction, meaning that the opposite contract can be bought at exactly $1 - \pi$ and that nothing is charged on either trade beyond the price. The first assumption is what makes prediction markets work. The second is false in every real market.

The most basic friction is the bid-ask spread. A market has a *bid*, the highest price anyone is currently willing to pay for Yes, and an *ask*, the lowest price at which anyone is currently willing to sell Yes. Buying is done at the ask and selling at the bid. Because selling Yes is the same as buying No, the No contract is offered at $1 - \text{Bid}$, not at $1 - \text{Ask}$. Apply the theorem to Yes at the ask and to No at its ask of $1 - \text{Bid}$. The first condition gives $p \leq \text{Ask}$ and the second gives $1 - p \leq 1 - \text{Bid}$, that is, $p \geq \text{Bid}$. So in an efficient market with a spread but no other friction,

$$\text{Bid} \leq p \leq \text{Ask}. \tag{1}$$

Suppose a contract has a bid of \$0.49 and an ask of \$0.50. A trader who believes the probability is 49.5% has no profitable trade. Buying at \$0.50 is a half-cent overpay, and selling at \$0.49 is a half-cent underpay. The same is true for any probability between 49% and 50%. The market does not say that the probability is 49% or that it is 50%. It says that the probability is somewhere between the two, and it is silent about where. Even with no

fees at all, the theorem turns the price into a range rather than a single number. Fees and commissions widen the range further.

3 The Inefficiency Interval and the Margin of Inefficiency

A *fee* is charged regardless of whether the contract wins or loses. A fee is effectively an increase in the price paid for the contract. A *commission* is charged only on winning contracts, usually as a percentage of profit. A commission is effectively a decrease in the amount the contract pays out when it wins. Putting these together, every contract has an *effective price*, which is the quoted price plus any fees, and an *effective settlement value*, which is \$1.00 less any commission. A trade is profitable in expectation only if

$$\text{probability} \times \text{effective settlement value} > \text{effective price}.$$

Apply this inequality to the Yes contract and the probability of Yes can be no larger than the effective price of Yes divided by its effective settlement value. Otherwise buying Yes would be profitable. Apply it to the No contract and the probability of No can be no larger than the effective price of No divided by its effective settlement value, which gives a lower limit on the probability of Yes. Together,

$$1 - \frac{\text{effective price of No}}{\text{effective settlement of No}} \leq \text{Pr(Yes)} \leq \frac{\text{effective price of Yes}}{\text{effective settlement of Yes}}. \quad (2)$$

The **inefficiency interval** of a market is the range of probabilities in (2). It is the set of probabilities for which no trade at the current prices is profitable, after all fees and commissions are paid. Every probability in the interval is consistent with the current market price.

The **margin of inefficiency** (MOI) is the width of the inefficiency interval.

The interval is computed from four numbers: the effective price and the effective settlement value of the Yes contract, and the same two numbers for the No contract. All four can be read off the order book and the platform’s fee schedule. No assumption about the probability is needed.

California, 2020. Return to the Biden contract on PredictIt from the introduction. On election morning the contract had a bid of \$0.93 and an ask of \$0.94. PredictIt charged a 10% commission on profits, and we treat the 5% withdrawal fee as a flat fee of \$0.05 per contract. Buying Yes at \$0.94 has an effective price of \$0.99. If it wins, the \$0.06 profit is charged \$0.006 in commission, so the effective settlement value is \$0.994. The upper limit of the interval is $0.99/0.994 = 99.6\%$. On the other side, buying No means paying \$0.07 (the complement of the \$0.93 bid) plus the \$0.05 fee, for an effective price of \$0.12, and a winning No contract is charged 10% of its \$0.93 profit, for an effective settlement value of \$0.907. The lower limit is $1 - 0.12/0.907 = 86.8\%$.

So the PredictIt price of \$0.94 placed Biden’s probability of winning California anywhere between 86.8% and 99.6%. The margin of inefficiency is 12.8 points. On Betfair, which charged a 5% commission and no fee, the same contract had a bid of \$0.98 and an ask of \$0.99. Its inefficiency interval runs from 97.9% to 99.1%, a margin of 1.2 points. The Betfair interval sits entirely inside the PredictIt interval. The two markets were not in disagreement about California. One of them simply said much less than the other.

The 2026 fee schedules. Neither of the two largest U.S. platforms for political markets charges a commission. Both charge a fee on each contract, paid by the trader who takes an existing order, and the size of the fee depends on the price. Kalshi charges $0.07 \times P \times (1 - P)$ per contract at price P [Kalshi, 2026]. Polymarket, which charged nothing in 2024, now charges $0.04 \times P \times (1 - P)$ per contract on political markets [Polymarket, 2026]. At a price of \$0.50 these fees are 1.75 cents and 1.00 cent per contract, and they shrink toward zero as the price approaches \$0.00 or \$1.00. Under a fee of this form at rate r , the inefficiency interval is

$$\text{Bid} - r \text{ Bid} (1 - \text{Bid}) \leq \text{Pr}(\text{Yes}) \leq \text{Ask} + r \text{ Ask} (1 - \text{Ask}), \quad (3)$$

where Bid and Ask are the bid and ask on Yes.

Table 1 shows what the same quote, a bid of \$0.49 and an ask of \$0.50, says about the probability under four different fee structures. The quote is identical in every row. Its meaning is not.

Fee structure	Inefficiency interval	MOI (points)
None	49.0% to 50.0%	1.0
Polymarket, 2026 ($0.04 \times P(1 - P)$ per contract)	48.0% to 51.0%	3.0
Betfair (5% commission on profit)	47.7% to 51.3%	3.6
Kalshi ($0.07 \times P(1 - P)$ per contract)	47.2% to 51.8%	4.5

Table 1: What a bid of \$0.49 and an ask of \$0.50 says about the probability of Yes, under four fee structures.

A useful rule of thumb is that each source of friction adds its own amount to the margin. For a fee of f per contract and a commission of κ on profit, with P the midpoint of the bid and the ask,

$$\text{MOI} \approx \text{bid-ask spread} + 2f + 2\kappa P(1 - P). \quad (4)$$

The fee counts twice because it is charged on both the Yes and the No side. For the price-dependent fees on Kalshi and Polymarket, the fee term is $2r P(1 - P)$. Figure 1 plots the margin at every price for the two platforms, assuming a one-cent spread. The margin is largest at \$0.50, which is exactly where the question of who is favored attracts the most attention, and it narrows to the spread as the price approaches \$0.00 or \$1.00.




Figure 1: Margin of inefficiency by price under the 2026 fee schedules on Kalshi and Polymarket (political markets), with a bid-ask spread of \$0.01.

4 Sources of Inefficiency

Though efficiency is the ideal for a market, inefficiency is the norm. And for analytical purposes, inefficiency can be measured while efficiency cannot. The margin of inefficiency measures the part of a market's inefficiency that is built into its design. What that includes, and what it leaves out, is as follows.

The bid-ask spread, fees, and commissions are *endogenous* sources of inefficiency. They are set by the platform, they are published or visible in the order book, and they apply equally to every trader. Formula (2) accounts for them exactly. Minimum tick size belongs in this category as well, since the spread can never be smaller than one tick. A one-cent tick is of little consequence at \$0.50, but at \$0.02 the same tick means that the market cannot distinguish a 1% probability from a 2% probability, which differ by a factor of two.

Other sources are *exogenous*. They are real, but they cannot be read off a fee schedule, and they differ from one trader to the next. Capital committed to a contract is tied up until the contract settles, which for political markets can be months or years, and the same capital could be earning interest or deployed on other opportunities in the meantime. A position in a thinly traded market may be difficult to exit before settlement. Position limits and participation restrictions, such as PredictIt's \$850 cap, limit the capital available to correct a price even when the correction is obviously profitable. And traders rarely know the true probability with certainty, so they typically require some margin of safety before trading. A trader who believes the probability is 51% will usually not buy at \$0.50, because a one-point edge leaves no room for error.

Exogenous sources are handled by adding an allowance to each end of the interval. The allowance can be based on the cost of capital. At an annual rate of 6%, for example, holding

a contract for two months costs about 1% of its price, which widens the interval by about one point on each side. The allowance can also be estimated from data, by taking several markets on the same event and asking how much each interval has to be widened before all of them overlap. For liquid markets on the major platforms within a few months of settlement, the data from the 2024 U.S. elections put this allowance between one and one and a half points. For thinly traded markets it is several times larger [Crane, 2026].

5 Consensus Across Markets

In practice, the true probability is never known, so a single market cannot be declared efficient or inefficient on its own. What can be done is to compare markets on the same event. If each market’s price is consistent with the true probability, then the true probability lies in every market’s inefficiency interval, and therefore in the overlap of all of them.

The **consensus** of several markets on the same event is the overlap of their inefficiency intervals. It runs from the largest of the lower limits to the smallest of the upper limits. If the overlap is non-empty, the markets are **coherent**. If it is empty, they are **incoherent**.

Coherent markets are not in disagreement, no matter how different their prices look. Their prices differ by no more than their frictions can explain, in the same sense that two polls with different results can both be within the margin of error. When the intervals fail to overlap, the gap between them is the amount of the price difference that fees and spreads cannot account for. The gap must be explained by something else: an exogenous friction, a difference in how the two contracts are worded or settled, or a genuine disagreement between the two markets that the available capital has not yet corrected.

Note that coherence says nothing about which market is better. Quite the contrary. A market with heavy fees and little liquidity will almost always be coherent with a liquid, low-fee market, because its interval is wide enough to contain almost any price. The consensus is driven by the narrowest intervals, which belong to the markets with the lowest friction. A market with a wide margin contributes little to the consensus, and its price should be given correspondingly little weight.

The 2024 presidential election. Table 2 shows the price for Donald Trump to win the 2024 election on five platforms one month before the election. The prices themselves span less than a point. The margins of inefficiency range from less than a point on Polymarket, which charged no fees at the time, to more than seventeen points on PredictIt. All five intervals overlap, and the consensus is 49.3% to 49.8%. That consensus is essentially Polymarket’s interval, because Polymarket’s interval sits inside every other. PredictIt’s price, by contrast, carried almost no information beyond what the other four platforms already said.

Platform	Trump price	Inefficiency interval	MOI (points)
Polymarket	0.496	49.3% to 49.9%	0.6
Smarmets	0.488	47.8% to 49.8%	2.0
Betfair	0.487	46.9% to 50.5%	3.6
Kalshi	0.495	47.0% to 52.0%	5.0
PredictIt	0.490	40.3% to 57.7%	17.3
Consensus		49.3% to 49.8%	0.5

Table 2: Prices and inefficiency intervals for Donald Trump to win the 2024 U.S. presidential election on October 4, 2024, including an allowance of half a point for one month of capital cost. From Crane [2026].

The seven-point gap between PredictIt and Polymarket described in the introduction opened later in October. On October 28, Trump was \$0.59 on PredictIt and \$0.66 on Polymarket. PredictIt’s inefficiency interval that day, based on its fees and spread alone, ran from 50.6% to 67.2%. Polymarket’s price was inside it. In fact, the PredictIt and Polymarket intervals overlapped on every one of the 32 days from October 4 through the eve of the election. The gap that was reported as evidence of partisan bias on one platform or manipulation on the other was, throughout, within PredictIt’s margin of inefficiency.

The same holds more broadly. Over the month before the 2024 election, Polymarket’s price for Trump averaged 2.5 points above Kalshi’s, and 2.2 of those points are accounted for by Kalshi’s fee and spread. With a one-point allowance for exogenous friction, Polymarket, Betfair, Smarmets, and Kalshi were coherent in 90% of the hours in that month. In 2020, the gap between PredictIt and Betfair on the national race was within PredictIt’s margin on every day before the election. And across 96 state and district contests that traded on both Kalshi and Polymarket in 2024, the same one-point allowance produced coherence 96% of the time in the state-level presidential markets and 62% of the time in the thinly traded House markets. Full details are in Crane [2026].

6 The 2026 Midterm Elections

The 2026 midterms are the first U.S. general election in which both major platforms charge fees that depend on price. In 2024, Polymarket’s interval was narrower than every other and sat inside them, so the consensus was simply Polymarket’s interval. In 2026 the two intervals are of comparable width, neither need contain the other, and the consensus can be narrower than either one.

Figure 2 shows the markets for party control of the Senate and the House on the morning of October 8, 2026, four weeks before the election. On the Senate, the contract for Democratic control was bid \$0.62 and offered at \$0.63 on Kalshi, and bid \$0.63 and offered at \$0.64 on Polymarket. Taken as probabilities, the platforms differ by a point. Kalshi’s inefficiency interval is 60.4% to 64.6%, Polymarket’s is 62.1% to 64.9%, and the consensus is 62.1% to 64.6%. On the House, the consensus is 89.6% to 90.7%. In both cases the platforms are

coherent on fees and spread alone, with no exogenous allowance needed.




Figure 2: Inefficiency intervals and consensus for Democratic control of the Senate and the House, from quotes on Kalshi and Polymarket at about 10:30 a.m. Eastern time on October 8, 2026, based on fees and spread only.

The election portal that this document accompanies (<https://harrycrane.com/moi.html>) reports these quantities continuously as prices change. For each contest it displays the inefficiency interval on each platform, the consensus, and whether the platforms are coherent. The portal covers control of each chamber and the individual Senate, House, and gubernatorial races that trade on more than one platform.

7 What the Margin of Inefficiency Does and Does Not Say

The inefficiency interval is not a confidence interval. It makes no probabilistic statement about where the true probability is likely to be. It is the set of probabilities that the market's own prices cannot rule out. The correct reading of any market is that one of two things is true: either the true probability is within the interval, or it is outside the interval and the traders who could profit from correcting the price have not done so. The second possibility is what people mean when they call a market irrational or manipulated. The first is the default, and it should be ruled out before the second is entertained.

The margin of inefficiency is a measure of precision, not accuracy. A narrow margin means that the price pins the probability down closely. It does not mean that the probability is right, any more than a poll with a small margin of error is guaranteed to be unbiased. A wide margin means the price says little, regardless of how carefully the market's participants

trade. Because the endogenous margin depends only on the platform’s design and the state of its order book, it can be computed for any market at any time, compared across platforms, and tracked over the life of a contract. Anything that lowers friction narrows the margin. Anything that raises it widens the margin.

When reporting a single number from a market, the midpoint of the inefficiency interval is the natural point estimate, with the margin of inefficiency reported alongside it in the same way that a poll result is reported with its margin of error. The midpoint need not equal the last traded price or the midpoint of the bid and ask, especially on platforms with commissions, which pull the interval asymmetrically.

Two cautions apply when comparing markets. First, the contracts must settle on the same outcome under the same rules. Differences in wording or settlement are a legitimate source of price differences that the margin of inefficiency does not address. Second, the calculation requires live bids and asks. A last traded price records what one trader paid at some point in the past. It does not say what is available now, and a margin computed from it understates the true margin.

References

- Crane, H. (forthcoming). *Prediction Markets*. Book manuscript.
- Crane, H. (2026). Margin of inefficiency in prediction markets. Working paper.
- Kalshi (2026). Fee schedule. <https://kalshi.com/docs/kalshi-fee-schedule.pdf>. Accessed September 2026.
- Polymarket (2026). Trading fees. <https://help.polymarket.com/en/articles/13364478-trading-fees>. Accessed October 2026.