

Towards a Theory of Confidence in Market-Based Predictions

Rupert Freeman

Darden School of Business, University of Virginia, USA

FREEMANR@DARDEN.VIRGINIA.EDU

David M. Pennock

Department of Computer Science, Rutgers University, USA

DPENNOCK@DIMACS.RUTGERS.EDU

Daniel M. Reeves

Beeminder

DREEVES@BEEMINDER.COM

David Rothschild

Microsoft Research New York City, USA

DAVIDMR@MICROSOFT.COM

Bo Waggoner

Department of Computer Science, University of Colorado, USA

BWAG@COLORADO.EDU

Abstract

Prediction markets are a way to yield probabilistic predictions about future events, theoretically incorporating all available information. In this paper, we focus on the *confidence* that we should place in the prediction of a market. When should we believe that the market probability meaningfully reflects underlying uncertainty, and when should we not? We discuss two notions of confidence. The first is based on the expected profit that a trader could make from correcting the market if it were wrong, and the second is based on expected market volatility in the future. Our paper is a stepping stone to future work in this area, and we conclude by discussing some key challenges.

Keywords: Prediction Markets, Forecasting

1. Introduction

A *prediction market* is a financial market directly designed to crowdsource predictions. Bettors trade securities that eventually realize some value based on the outcome of the event in question. The price of the securities at any point in time can be interpreted as a forecast probability that the event occurs.

Given a prediction market forecast, it is natural to ask to what extent it reflects the actual uncertainty surrounding the event in question. If two markets selling the exact same contract disagree, which should we trust more? How confident can we be in the accuracy of the (aggregated or separate) predictions? Even though a prediction market has no concept of margin of error, can we design an instrument that properly conveys the accuracy of, and confidence in, a market prediction? In this paper, we attempt to paint a way forward to answering these questions.

By way of motivation, consider three scenarios predicting the winner of a football match:

1. a normal professional regular-season game

2. a normal professional regular-season game where, 24 hours before the game, there will be a rapid COVID-19 test to determine the eligibility of the star player
3. a pick-up game

Now imagine that a crowdsourced forecast (via whatever method) gives Team A a 70% chance of winning 48 hours before the game in all three scenarios. Intuitively, that crowd knows that Game 1 is very precisely defined due to years of detailed statistics and precise models. In Game 2, we will treat it as known that the star player's COVID-19 test is essentially a fair coin flip and that Team A has a 50% chance of winning without their star and 90% with them. That gives a multi-modal future probability [14]: $0.5 \cdot 90\% + 0.5 \cdot 50\% = 70\%$. Finally, the Game 3 estimate is based on scant information; who knows who may show up to play or what may happen? For all three games, 70% for Team A is our probability, but that 70% number is distilling intuitively very different states of uncertainty about the outcomes.

In this paper, we ask whether it is possible to formalize that feeling of confidence: to ascertain and explain which, if any, of the three football matches is most reminiscent of a given real-world probabilistic prediction. That is, we seek to develop a theory of confidence in predictions. While we feel that this is an intriguing topic in any setting where predictions are made, we will focus our attention on prediction markets, given their widespread use and influence in popular forecasting domains such as sports and politics.

Related Work A large body of literature focuses on the relative accuracy of prediction markets and other wisdom-of-crowds forecasts compared to polls and traditional expert forecasting methods [18, 19, 6, 9, 15]. However, the majority of these papers do not formally examine the notion of confidence, or consider ex-ante forecast accuracy. Closest to the notion of market confidence, Berg et al. [3, 4] empirically examine prediction market accuracy using data

from the Iowa Electronic Markets. However, they focus on vote share markets as opposed to probabilistic predictions. Frongillo et al. [7] study the problem of aggregating probabilistic beliefs when forecasters may have differing levels of information, but assumes that information is provided directly to the principal and not via a market.

Work on imprecise probabilities [1] has studied the problem of aggregating multiple imprecise probabilities [13, 17, 12], generally from an axiomatic perspective. Other work has considered the meaning and necessity of a notion of second-order probabilities to express uncertainty over a probability [2, 10, 11, 14]. These are relevant literatures for the direction we propose, although to our knowledge, existing work does not directly shed light on market confidence.

2. A Confidence Measure from the Efficient Market Hypothesis

Why do we trust that (prediction) market prices reflect the underlying value of an asset? The *efficient market hypothesis* states that markets incorporate all available information. The idea behind the hypothesis is that if some relevant piece of information is not reflected in the market price, then someone with that piece of information could profit from buying or selling the asset until its price reaches the appropriate level. Since the existence of such profit opportunities is not a stable state of a market, in equilibrium we would expect that all information is incorporated.

Unfortunately, real-world markets may not be fully efficient, with information often failing to flow into the market. In these cases, opportunities to profit can persist — indeed, many people make careers out of finding and exploiting these opportunities. But we can at least expect small profit opportunities to persist longer and more reliably than large ones. We expect a \$20 bill to go unnoticed on the sidewalk longer than a pot of gold!

In this section we exploit the ideas behind the efficient market hypothesis to propose a general method for expressing confidence in a market prediction. To set the stage more formally, suppose that we observe a prediction market for a random bit of uncertainty to be realized at some known time in the future. How can we estimate the probability p that the random bit takes value 1, and how can we express our confidence in that estimate?

Consider a prediction market with current price q . In our football match example, we have $q = 70\%$ as the market estimate. Now imagine that q is not actually representative of p , the underlying randomness in the event. Say $q = p + \varepsilon$ for some $\varepsilon > 0$. Then an omniscient being with knowledge of p would be able to make some expected profit $\$x > 0$ by moving the market price to p , exploiting all available trade in the market by doing so. The existence of this undiscovered profit opportunity would be at least mildly surprising, certainly more surprising than $q = p$ and all profit opportu-

nities having been exploited. In this sense we can quantify exactly how surprised we would be if $p = q + \varepsilon$: we would be $\$x$ surprised. The greater the ε , the greater the market inefficiency, and the greater the surprise.

Real-world prediction markets incorporate trading fees, may have low liquidity, or cap the amount that a single trader can invest. Tying up capital also entails opportunity cost. All of this distorts our ability to infer a probability from the market. For example, suppose that a market platform charges a flat 5% fee on the amount paid in any trade. Then, to make a profit, a trader must be able to buy or sell securities at a price at least 5% removed from what she estimates to be their value. In such a market, we could only expect the discovery of an accurate probability up to a 5% margin of error. In other words, if the true probability p was anywhere within 5% of the market price q , the omniscient profit would be \$0, just the same as if $p = q$ exactly.

This discussion therefore yields a natural way to discuss uncertainty in probabilities arising from prediction markets. For any market, and any probability p , we can provide an exact amount of money that quantifies how surprised we would be if the true probability took value p . This method is very general. We simply put ourselves in the shoes of a (budget unconstrained) omniscient trader who knows that p is the true probability that the event occurs. Then we imagine participating in the market, exhausting all possible trades that yield a positive expected profit, net of fees. The omniscient’s expected profit $\$x$ is our level of surprise, should the true probability indeed be p .

Conversely, for an amount of money $\$x$, we can consider the set of all possible true probabilities that would yield an omniscient profit of at most $\$x$. In this way we can define an “ $\$x$ margin of error” for any prediction market, as the set of probabilities that are consistent with the current state of the market up to an $\$x$ level of surprise.

Note that many features of a market that we intuitively associate with increased accuracy will also be associated with narrow margins of error: high liquidity, no investment caps, and low fees. All of these features increase the (expected) profit that an omniscient trader could extract from participating in the market, shrinking the space of true probabilities that would result in a profit less than $\$x$.

Consider (hypothetical) prediction markets for the professional and pickup football matches described earlier. Due to the high information flow and level of interest in the professional match, we would expect a thick market with many opportunities for trade. The margin of error in this market will be narrow, because a trader with knowledge that the market was wrong could make a large profit. On the other hand, we would expect the pickup game market to be thinly trafficked. Even a trader who had precise knowledge of the capabilities of the teams would be able to make very little money. The margin of error in this market would be wide, reflecting a low confidence in the market estimate.

One can imagine prediction market estimates being reported in this way. For example, a data journalist could report that the probability of Joe Biden being elected president is 65%, with the \$1000 margin of error being between 57% and 70%. A market with a narrower margin of error at the $\$x$ level should be more trustworthy than a market with a wider margin of error.

This method can even extend to settings with multiple markets for the same event, or situations with very general betting structures that need not take the form of conventional prediction markets. When there are multiple markets (say, two different platforms each having a market for the same event), we can imagine an informed trader who is able to invest in all markets, profiting from each one. Since a trader can make higher expected profit from participating in two markets than just one, the margin of error that results from considering multiple markets will be no wider than the margin of error from considering a subset of those markets. Once again, this matches our intuition; additional information should not make us *less* confident in our prediction. Similarly, in any betting situation, we can ask how much profit a bettor with infinite budget and precise knowledge of the underlying true probability p could make in expectation, given the bets laid down by the other parties.

3. Confidence as Volatility

Let us return to Game 2 from the introduction, in which a 50/50 random event will occur *before* the football match in question, the outcome of which will inform us whether Team A has a 50% or a 90% chance of winning the match.

Can we say we're confident in today's estimate of 70%? By the measure we suggest in Section 2, yes. The information structure is public so we expect plenty of traders willing to bet if the implied odds deviate from 70% in either direction. But in another sense, no. We're expecting new information, 24 hours before the game, that will render the 70% estimate wrong in hindsight. A layperson told that a probabilistic prediction has high confidence would be taken aback to see the prediction change drastically the following day. We can capture that notion of confidence (or lack thereof) by measuring the expected *volatility*. A prediction with high expected volatility indicates, in our toy example, additional layers of uncertainty that will be resolved before the event itself resolves.

In financial markets, options and other derivatives reveal the distribution and variance of their underlying instruments. For example, a *butterfly option* pays the absolute difference between the market price at time t_1 and the price at time t_2 : $|p_{t_1} - p_{t_2}|$. Such contracts could be traded in a market secondary to any prediction, allowing us to estimate the volatility of the primary prediction.

In other words, a *second-order* market can predict volatility in the original market. To be clear, second-order probabilities can always be collapsed into first-order probabili-

ties [14]. Nonetheless, the concept of second-order probabilities are well-defined and capture the expectation of new information that will change our probability estimate. For the example of Game 2, the price of the butterfly option would reveal that the current prediction, a 70% chance of Team A winning, is expected to change by $\pm 20\%$ exactly when the COVID-19 test occurs. This would reflect a lack of confidence in the prediction, at least relative to that point in the future.

We believe that this “volatility” notion of confidence is complementary to the “ $\$x$ margin of error” notion of confidence. One can imagine scenarios, such as the Game 2 example, where the $\$x$ margin of error is very small while the volatility is large. The converse appears possible, but less likely: a high $\$x$ margin of error implies high uncertainty about the current prediction, which suggests volatility of the price in the future. In any case, both measures together may give a clearer picture of the market's confidence.

4. Discussion and Challenges

We have proposed two ways to measure the confidence of a prediction market forecast; one appealing to the efficient market hypothesis and the other concerned with the expected volatility in the market. However, before these notions can be used, they need to be validated empirically and/or theoretically. A natural step would be to check whether high-confidence predictions are in fact correlated with accuracy, as measured by a proper scoring rule such as the Brier score [5, 8] that rewards accurate forecasts more than inaccurate ones in expectation.

The notions of confidence in Sections 2 and 3 can both be interpreted in terms of the value of information acquisition. In Section 2, an equivalent interpretation of the $\$x$ margin of error is that it contains all the predictions that could result if a forecaster spent $\$x$ acquiring new information (otherwise, a forecaster could make greater than $\$x$ profit from the market by spending only $\$x$ on information acquisition, violating the efficient market hypothesis¹). In Section 3, the expected market volatility expresses how much we expect the prediction to change as new information comes to light. High expected volatility says that the new information will be highly valuable compared to old information. Further exploring the relationship between value of information and market confidence may be fruitful.

Finally, much of our discussion has implicitly assumed that every event has some inherent and unknowable, yet quantifiable and well-defined, uncertainty occurring at a precise moment in time. Alternative models are of course possible, and may affect how we define and think about uncertainty. The interplay between these ideas and other game-theoretic notions of probability [16] may also be fruitful.

1. In reality, we don't expect markets to be so efficient that this holds exactly.

Acknowledgments

We are grateful to Jennifer Wortman Vaughan for helpful discussions and to Bethany Soule for feedback on a preliminary version of the paper.

References

- [1] Thomas Augustin, Frank PA Coolen, Gert De Cooman, and Matthias CM Troffaes. *Introduction to Imprecise Probabilities*. John Wiley & Sons, 2014.
- [2] Jonathan Baron. Second-order probabilities and belief functions. *Theory and Decision*, 23(1):25–36, 1987.
- [3] Joyce Berg, Robert Forsythe, and Thomas Rietz. What makes markets predict well? Evidence from the Iowa Electronic Markets. In *Understanding Strategic Interaction*, pages 444–463. Springer, 1997.
- [4] Joyce Berg, Forrest Nelson, and Thomas Rietz. Accuracy and forecast standard error of prediction markets. *Working paper*, 2003.
- [5] Glenn W Brier. Verification of forecasts expressed in terms of probability. *Monthly Weather Review*, 78(1):1–3, 1950.
- [6] Robert S Erikson and Christopher Wlezien. Are political markets really superior to polls as election predictors? *Public Opinion Quarterly*, 72(2):190–215, 2008.
- [7] Rafael M Frongillo, Yiling Chen, and Ian A Kash. Elicitation for aggregation. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, pages 900–906, 2015.
- [8] Tilmann Gneiting and Adrian E Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American statistical Association*, 102(477):359–378, 2007.
- [9] Sharad Goel, Daniel M Reeves, Duncan J Watts, and David M Pennock. Prediction without markets. In *Proceedings of the 11th ACM conference on Electronic Commerce*, pages 357–366, 2010.
- [10] Robert W Goldsmith and Nils-Eric Sahlin. The role of second-order probabilities in decision making. In *Advances in Psychology*, volume 14, pages 455–467. Elsevier, 1983.
- [11] Sven Ove Hansson. Do we need second-order probabilities? *Dialectica*, 62(4):525–533, 2008.
- [12] Serafín Moral and Jose Del Sagrado. Aggregation of imprecise probabilities. In *Aggregation and Fusion of Imperfect Information*, pages 162–188. Springer, 1998.
- [13] Robert F Nau. The aggregation of imprecise probabilities. *Journal of Statistical Planning and Inference*, 105(1):265–282, 2002.
- [14] Judea Pearl. Do we need higher-order probabilities and, if so, what do they mean? In *Proceedings of the Third Conference on Uncertainty in Artificial Intelligence*, pages 47–60, 1987.
- [15] David Rothschild. Combining forecasts for elections: Accurate, relevant, and timely. *International Journal of Forecasting*, 31(3):952–964, 2015.
- [16] Glenn Shafer and Vladimir Vovk. *Game-Theoretic Foundations for Probability and Finance*, volume 455. John Wiley & Sons, 2019.
- [17] Rush T Stewart and Ignacio Ojea Quintana. Probabilistic opinion pooling with imprecise probabilities. *Journal of Philosophical Logic*, 47(1):17–45, 2018.
- [18] Justin Wolfers and Eric Zitzewitz. Prediction markets. *Journal of Economic Perspectives*, 18(2):107–126, 2004.
- [19] Justin Wolfers and Eric Zitzewitz. Prediction markets in theory and practice. Technical report, National Bureau of Economic Research, 2006.