Learning Performance of Prediction Markets with Kelly Bettors

(Extended Abstract)

Alina Beygelzimer
IBM Research
beygel@us.ibm.com

John Langford
Yahoo! Research
jl@yahoo-inc.com

David M. Pennock
Yahoo! Research
pennockd@yahoo-inc.com

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Intelligent agents, Multiagent systems*

General Terms

Economics

Keywords

Auction and mechanism design, electronic markets, economically motivated agents, multiagent learning

1. INTRODUCTION

Consider a prediction market, in which participants can trade *shares* (binary options) at the current market price p_m. Each share is worth \$1 if the event occurs, and nothing otherwise. What fraction of your wealth w should you risk if you believe the probability of the event is p? Buying is favorable if $p > p_m$, in which case risking your entire wealth will maximize your expected profit with respect to your belief. However, that's extraordinarily risky: A single stroke of bad luck loses everything. On the other hand, risking a small fixed amount cannot take advantage of compounding growth.

The *Kelly criteria* prescribes investing f^*w dollars, where for $p > p_m$,

$$f^* = \frac{p - p_m}{1 - p_m}$$

(buy order). For $p < p_m$, you should bet against the outcome (sell order) with

$$f^* = \frac{(1-p) - (1-p_m)}{1 - (1-p_m)} = \frac{p_m - p}{p_m}.$$

Kelly betting maximizes the expected compounding growth rate of wealth, or equivalently the expected logarithm of wealth [2, 4, 10].

We consider a prediction market, where participant i starts with wealth w_i, with $\sum_i w_i = 1$. Each participant i uses Kelly betting to determine the fraction of their wealth to bet, depending on their prediction p_i.

Appears in: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, Conitzer, Winikoff, Padgham, and van der Hoek (eds.), 4-8 June 2012, Valencia, Spain.

We model the market as an auctioneer matching supply and demand, taking no profit and absorbing no loss, with p_m selected to clear the market. Agents are "price takers" that optimize according to the current price and do not reason further about what the price might reveal about the other agents' information. (In the fractional Kelly setting, however, agents do consider the market price as information and weigh it along with their own.)

Observation 1. *The market prediction p_m is always a wealth-weighted average of the agents' predictions p_i,*

$$p_m = \sum_i p_i w_i.$$

Proof. The market equilibrium occurs at price p_m where the payin is equal to the payout. If the event occurs,

$$\sum_{i:p_i>p_m} \frac{p_i - p_m}{1-p_m} w_i + \sum_{i:p_i<p_m} \frac{p_m - p_i}{p_m} w_i = \frac{1}{p_m} \sum_{i:p_i>p_m} \frac{p_i - p_m}{1-p_m} w_i.$$

Simplifying, we get $\sum_i p_i w_i = p_m \sum_i w_i$. Applying $\sum_i w_i = 1$ finishes the proof. A similar calculation proves the observation if the event doesn't occur. □

An alternate derivation utilizes the fact that Kelly betting is equivalent to maximizing expected log utility. This result can be seen as a simplified derivation of that by Rubinstein [7, 8, 9] and is also discussed by Pennock and Wellman [6, 5] and Wolfers and Zitzewitz [11].

2. LEARNING PREDICTION MARKETS

Consider a sequence of prediction markets which may have varying true and predicted probabilities. What happens to the wealth distribution and hence the quality of the market prediction over time? We show that the market *learns* optimally for two well understood senses of optimal.

2.1 Wealth is redistributed according to Bayes' Law

In an individual round, if an agent's belief is $p_i > p_m$, their total wealth afterward depends on the outcome y according to

$$\text{If} \quad y=1, \quad \left(\frac{1}{p_m} - 1\right)\frac{p_i - p_m}{1-p_m} w_i + w_i = \frac{p_i}{p_m} w_i$$

$$\text{If} \quad y=0, \quad -\frac{p_i - p_m}{1-p_m} w_i + w_i = \frac{1-p_i}{1-p_m} w_i$$

Similarly if $p_i < p_m$, we get

$$\text{If} \quad y = 1, \quad -\frac{p_m - p_i}{p_m} w_i + w_i = \frac{p_i}{p_m} w_i$$

$$\text{If} \quad y = 0, \quad \left(\frac{1}{1-p_m} - 1\right) \frac{p_m - p_i}{p_m} w_i + w_i = \frac{1-p_i}{1-p_m} w_i,$$

which is identical.

If we treat the prior probability $P(i)$ that agent i is correct as w_i, the posterior probability of choosing agent i is

$$P(i \mid y = 1) = \frac{P(y = 1 \mid i)P(i)}{\sum_j P(y = 1 \mid j)P(j)} = \frac{p_i w_i}{p_m},$$

which is precisely the wealth computed above for the $y = 1$ outcome, and similarly when $y = 0$. So Kelly bettors redistribute wealth according to Bayes' law, and the market price reacts exactly as if updating according to Bayes' law.

In the full version [1], we simulate a sequence of markets where an underlying true probability exists, showing that the market converges to the true objective frequency as if updating a Beta distribution, as the theory predicts.

Although individual agents are not adaptive, the market's composite agent computes a proper Bayesian update. Specifically, wealth is reallocated proportionally to a Beta distribution corresponding to the observed number of successes and trials, and price is approximately the expected value of this Beta distribution. A kind of collective Bayesianity emerges from the interactions of the group.

We also find empirically that, even if not all agents are Kelly bettors, among those that are, wealth is still redistributed according to Bayes' rule.

2.2 Market has low regret to the best agent

The assumptions in the section above are often too strong. The following result applies to *all* sequences of participant predictions p_{it} and *all* outcome sequences y_t, even when these are chosen adversarially. It states that even in this worst-case situation, the market performs no worse than $-\ln w_i$ compared to the best individual participant i, using standard analysis from learning theory [3].

We measure the accuracy of market predictions $\{p_t\}$ according to log loss as

$$L \doteq \sum_{t=1}^{T} I(y_t = 1) \log \frac{1}{p_t} + I(y_t = 0) \log \frac{1}{1-p_t}.$$

Similarly, the accuracy of participant i is measured as

$$L_i \doteq \sum_{t=1}^{T} I(y_t = 1) \log \frac{1}{p_{it}} + I(y_t = 0) \log \frac{1}{1-p_{it}}.$$

THEOREM 2. *For all sequences of participant predictions p_{it} and all sequences of revealed outcomes y_t,*

$$L \leq \min_i L_i + \ln \frac{1}{w_i}.$$

PROOF. Initially, we have $\sum_i w_i = 1$. After T rounds, the total wealth of any participant i is given by

$$w_i \prod_{t=1}^{T} \left(\frac{p_{it}}{p_t}\right)^{y_t} \left(\frac{1-p_{it}}{1-p_t}\right)^{1-y_t} = w_i e^{L-L_i} \leq 1,$$

where w_i is the starting wealth and the last inequality follows from wealth being conserved. Thus $\ln w_i + L - L_i \leq 0$, yielding $L \leq L_i + \ln \frac{1}{w_i}$. $\square$

Thus self-interested agents with log wealth utility create markets which learn to have small regret according to log loss.

3. FRACTIONAL KELLY BETTING

In the full version of the paper [1], we consider fractional Kelly betting, a commonly used, lower-risk variant of Kelly betting, and show that fractional Kelly agents behave like Kelly agents with beliefs weighted between their own and the market's. When a true underlying probability exists, the market price empirically converges to a time-discounted version of this probability [1]. We also propose a method for agents to learn their optimal fraction over time.

4. QUESTIONS

When agents have some utility other than log wealth utility, can we alter the structure of a market so that the market dynamics make the market price have low log loss regret? And similarly if we care about some other loss—such as squared loss, 0/1 loss, or a quantile loss—can we craft a marketplace such that log wealth utility agents achieve small regret with respect to these other losses?

5. REFERENCES

[1] A. Beygelzimer, J. Langford, and D. M. Pennock. Learning performance of prediction markets with Kelly bettors. *CoRR (arXiv)*, abs/1201.6655, 2012.

[2] L. Breiman. Optimal gambling systems for favorable games. In *Berkeley Symposium on Probability and Statistics, I*, pages 65–78, 1961.

[3] N. Cesa-Bianchi and G. Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, New York, NY, USA, 2006.

[4] J. Kelly. A new interpretation of information rate. *Bell System Technical Journal*, 35:917–926, 1956.

[5] D. M. Pennock. *Aggregating Probabilistic Beliefs: Market Mechanisms and Graphical Representations*. PhD thesis, University of Michigan, 1999.

[6] D. M. Pennock and M. P. Wellman. A market framework for pooling opinions. Technical Report 2001-081, NEC Research Institute, 2001.

[7] M. Rubinstein. An aggregation theorem for securities markets. *Journal of Financial Economics*, 1(3):225–244, 1974.

[8] M. Rubinstein. Securities market efficiency in an Arrow-Debreu economy. *Americian Economic Review*, 65(5):812–824, 1975.

[9] M. Rubinstein. The strong case for the generalized logarithmic utility model as the premier model of financial markets. *Journal of Finance*, 31(2):551–571, 1976.

[10] E. O. Thorp. Optimal gambling systems for favorable games. *Review of the International Statistical Institute*, 37:273–293, 1969.

[11] J. Wolfers and E. Zitzewitz. Interpreting prediction market prices as probabilities. Technical Report 12200, NBER, 2006.