When Do Prediction Market Makers Make Profits?

Po-An Chen

Institute of Information Management, National Yang Ming Chiao Tung University

This is a joint work with Yiling Chen, Chi-Jen Lu, and Chuang-Chieh Lin

August 16, 2021

Outline

- Introduction: Prediction Markets
- Preliminaries
  - Designing the Cost Function via Congugate Duality
  - Equivalence between Market Making and Online Learning
  - Predictable Sequences
- Profitable Market Making via Online Learning
  - Optimistic Lazy-Update Online Mirror Descents with Predictable Sequences
  - Expert Setting: Modified Double Multiplicative Updates with Predictable Sequences
- Conclusions and Future Work

Prediction Markets

- Arrow-Debreu securities [2]: payoffs depend on the future state of the world.
    - Certain security pays 1 if a particular state of the world is reached.
    - A risk-neutral trader believing that the prob. of event occurring is p
    - Buying/selling this security if the price is below/above p
    - E.g., insurance contracts, options, futures, and many other financial derivatives
- The information of the trader can be capitalized.
- The market price of the security: traders' collective estimate of how likely some event will occur
- *Securities markets*: mechanisms for *risk allocation* and *information aggregation*
- Focusing on information aggregation: *prediction markets*

Market Makers

- An *automated market maker*: an institution adaptively setting prices for each security and accepting trades at these prices all the time
- A common goal of a prediction market making: *the worst-case loss*
- A securities market is *complete* if it offers at least $\mathcal{O}$ linearly independent securities over a set $\mathcal{O}$ of mutually exclusive and exhaustive outcomes [2].
  - A trader may bet on any combination of the securities in a complete securities market.
- **We aim to achieve market making guaranteeing profits, i.e., negative regrets, under appropriate patterns of trade sequences, requiring conditions other than just low deviation or variation.**

Designing the Cost Function via Congugate Duality

- Duality-based cost function market maker [1]
  - Input: outcome space $\mathcal{O}$
  - Input: K securities by a payoff function $\rho : \mathcal{O} \to \mathbb{R}^K_{\geq 0}$
  - Input: Convex compact price space Π (typically $\Pi = \mathcal{H}(\rho(\mathcal{O}))$)
  - Input: Strictly convex R with $relint(\Pi) \subseteq dom(R)$
  - Output: Market mechanism by the cost function $C : \mathbb{R}^K \to \mathbb{R}$ with

$$C(\mathbf{q}) := \sup_{\mathbf{x} \in relint(\Pi)} \mathbf{x} \cdot \mathbf{q}_{t-1} - R(\mathbf{x}) \tag{1}$$

- We can always write

$$C(\mathbf{q}) = \max_{\mathbf{x} \in \mathcal{H}(\rho(\mathcal{O}))} \mathbf{x} \cdot \mathbf{q} - R(\mathbf{x}) \tag{2}$$

Equiv. between Market Making and Online Learning (1/3)

- Lemma ([1]): If R is continuous and defined on all of $\mathcal{H}(\rho(\mathcal{O}))$, the price vector at any $\mathbf{q} \in \mathbb{R}^K$ satisfies

$$\nabla C(\mathbf{q}) = \arg \max_{\mathbf{x} \in \mathcal{H}(\rho(\mathcal{O}))} \mathbf{x} \cdot \mathbf{q} - R(\mathbf{x}) \tag{3}$$

- A duality-based cost function market maker:
  1. The market maker having an outcome space $\mathcal{O}$ and payoff function ρ, a feasible price space $\Pi = \mathcal{H}(\rho(\mathcal{O}))$.
  2. Select instantaneous security prices $x_t \in \Pi$.
  3. Uses a convex conjugate $R(\cdot)$, a parameter of the pricing function $C(\cdot)$.
  4. Receives security bundle purchases r_t. A quantity vector q_t updated according to $q_t = q_{t-1} + r_t$
  5. Sets prices via $x_t = \arg\max_{x \in \Pi} x \cdot q_{t-1} - R(x)$.
  6. The market maker's worst-case loss: $C(q_0) - C(q_T) + \max_{x \in \Pi} x \cdot q_T$

Equiv. between Market Making and Online Learning (2/3)

- $$C(\mathbf{q}_T) - C(\mathbf{q}_0) = \sum_{t=1}^{T} C(\mathbf{q}_t) - C(\mathbf{q}_{t-1}) \tag{4}$$
  $$\simeq \sum_{t=1}^{T} \nabla C(\mathbf{q}_{t-1}) \cdot (\mathbf{q}_t - \mathbf{q}_{t-1}) = \sum_{t=1}^{T} \mathbf{x}_t \cdot \mathbf{r}_t$$
- The difference between $C(\mathbf{q}_t) - C(\mathbf{q}_{t-1})$ and $\nabla C(\mathbf{q}_{t-1}) \cdot (\mathbf{q}_t - \mathbf{q}_{t-1})$: $D_C(\mathbf{q}_t, \mathbf{q}_{t-1})$; with conditions, $D_C(\mathbf{q}_t, \mathbf{q}_{t-1}) = D_R(\mathbf{x}_t, \mathbf{x}_{t+1})$
- The worst-case loss of the market maker:

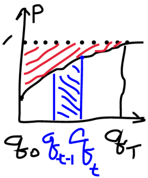

$\max_{\mathbf{x} \in \Pi} \mathbf{x} \cdot \mathbf{q}_T - \sum_{t=1}^{T} \mathbf{x}_t \cdot r_t - \sum_{t=1}^{T} D_R(\mathbf{x}_t, \mathbf{x}_{t+1})$

Equiv. between Market Making and Online Learning (3/3)

- Considering security bundles r_t as the negative loss vectors ℓ_t, $\longrightarrow$ **the duality-based cost function market maker becomes FTRL [3, 1].**
- Online linear optimization problem using FTRL:
  1. The learner given access to a fixed space of weights $\mathcal{K}$
  2. Selects a weight vector $x_t \in \mathcal{K}$.
  3. Uses a convex regularizer $\mathcal{R}(\cdot)$, a parameter of FTRL.
  4. Receives loss vectors f_t. A cumulative loss vector L_t updated according to $L_t = L_{t-1} + \ell_t$
  5. FTRL selects the weights by solving $x_t = \arg\min_{x\in\mathcal{K}}[\sum_{s=1}^{t-1} f_s(x) + \frac{1}{\eta}\mathcal{R}(x)]$.
  6. The learner's regret: $\sum_{t=1} x_t \cdot f_t - \min_{x\in\mathcal{K}} x \cdot L_T$

Predictable Sequences (1/2)

- "Regular" or "predictable" sequences [6]: a sequence of functions $M_t : \mathcal{F}^{t-1} \to \mathcal{F}$ for each $t \in \{1, ..., T\}$. A predictable process of the environment:

$$M_1, M_2(\ell_1), ..., M_T(\ell_1, ..., \ell_{T-1})$$

$\longrightarrow M_t(M_1, ..., M_{t-1}) = \ell_t$ for all t: $\{\ell_t\}$ is a noiseless time series.

- Optimistic agile-update online mirror descent ([5, 6])
  1. Let $x_1 = \hat{x}_1 = y_1 = (\frac{1}{N}, ..., \frac{1}{N})^T$.
  2. **for** $t \in [T]$ **do**

     Receive f_t and compute $\ell_t = \nabla f_t(\hat{x}^t)$.

     Update

     $x_{t+1} = \arg_{x \in \mathcal{X}} \mathcal{B}^{\mathcal{R}}(x, y_{t+1})$ for $\nabla \mathcal{R}(y_{t+1}) = \nabla \mathcal{R}(y_t) - \eta_t \ell_t$,

     $\hat{x}_{t+1} = \arg_{\hat{x} \in \mathcal{X}} \mathcal{B}^{\mathcal{R}}(\hat{x}, \hat{y}_{t+1})$ for $\nabla \mathcal{R}(\hat{y}_{t+1}) = \nabla \mathcal{R}(x_{t+1}) - \eta_t M_{t+1}$.

     end for

Predictable Sequences (2/2)

- L_p-deviation for the loss functions:
  $D_p = \sum_{t=1}^{T} \max_{\mathbf{x}\in\mathcal{X}} \|\nabla f_t(\mathbf{x}) - \nabla f_{t-1}(\mathbf{x})\|_p^2$
- Low deviation implies $\ell_t \simeq \ell_{t-1}$ so $M_{t+1} = \ell_t$.
- Theorem ([5]): When the L_2-deviation is D_2, the regret of the algorithm is at most $O(\sqrt{D_2})$.
- Theorem ([5]): When the L_∞-deviation is D_∞, the regret of the algorithm is at most $O(\sqrt{D_\infty \ln N})$.

Optimistic Lazy-Update Online Mirror Descents as Be The Regularized Leader with Predictors M_{t+1} (1/2)

- Optimistic lazy-update online mirror descent
  1. Let $x_1 = \hat{x}_1 = y_1 = (\frac{1}{N}, ..., \frac{1}{N})^T$.
  2. **for** $t \in [T]$ **do**

     Receive f_t and compute $\ell_t = \nabla f_t(\hat{x}_t)$.

     Update

     $x_{t+1} = \arg_{x \in \mathcal{X}} \mathcal{B}^{\mathcal{R}}(x, y_{t+1})$ for $\nabla \mathcal{R}(y_{t+1}) = \nabla \mathcal{R}(y_t) - \eta_t \ell_t$,

     $\hat{x}_{t+1} = \arg_{\hat{x} \in \mathcal{X}} \mathcal{B}^{\mathcal{R}}(\hat{x}, \hat{y}_{t+1})$ for $\nabla \mathcal{R}(\hat{y}_{t+1}) = \nabla \mathcal{R}(y_{t+1}) - \eta_t M_{t+1}$.

     end for
- Double lazy-update online mirror descent: $M_{t+1} = \ell_t$

Optimistic Lazy-Update Online Mirror Descents as Be The Regularized Leader with Predictors M_{t+1} (2/2)

- Lemma: The first update of optimistic lazy-update online mirror descents is $\nabla R_t(y_{t+1}) = \nabla R_t(y_t) - \eta_t \ell_t$ so

$$x_{t+1} = \arg\min_{x\in\mathcal{K}} \mathcal{B}^{R_t}(x, y_{t+1}) = \arg\min_{x\in\mathcal{K}} (\eta(\sum_{s=1}^{t} \ell_s^T)x + R(x)),$$

and the second update $\nabla R(\hat{y}_{t+1}) = \nabla R(y_{t+1}) - \eta_t M_{t+1}$ so

$$\hat{x}_{t+1} = \arg\min_{x\in\mathcal{K}} \mathcal{B}^{R_t}(x, \hat{y}_{t+1}) = \arg\min_{x\in\mathcal{K}} (\eta(\sum_{s=1}^{t} \ell_s^T + M_{t+1}^T)x + R(x)).$$

With $M_{t+1} = \ell_t$,

$$\hat{x}_{t+1} = \arg\min_{x\in\mathcal{K}} \mathcal{B}^{R_t}(x, \hat{y}_{t+1}) = \arg\min_{x\in\mathcal{K}} (\eta(\sum_{s=1}^{t} \ell_s^T + \ell_t^T)x + R(x)).$$

Strong Leaders with Frequent Changes of Leaders

- Lemma: For all $t \in \{1, ..., T\}$, with $L_t(x) - L_t(x_{t+1}) \geq \delta$ for all every other $x \notin S_t = \arg\min_x L_t(x)$ over k changes of leaders,
$$\sum_{t=1}^{T} \ell_t(x_{t+1}) - L_T(x^*) \leq -\delta k + \frac{1}{\eta_t}(R(x^*) - R(x_0)).$$
- Theorem: The regret is
$\sum_{t=1}^{T}(\ell_t(\hat{x}_t) - \ell_t(x_{t+1})) + \sum_{t=1}^{T} \ell_t(x_{t+1}) - L_T(x^*)$
$$= O(\sqrt{\sum_{t=1}^{T} \|\ell_t\|_2 \|\ell_t - M_t\|_2}) - \delta k,$$
where $x_{t+1} = \arg\min_{x \in \mathcal{K}}(\sum_{s=1}^{t} \ell_s(x) + \frac{R(x)}{\eta_t})$ and
$$\eta_t = \frac{1}{\sqrt{\sum_{s=1}^{t-1} \|\ell_s\|_2 \|\ell_s - M_s\|_2}}.$$

Reward Maximizing Expert Setting: Modified Optimistic Multiplicative Updates with Predictable Sequences (1/2)

- A reward maximization model with experts of $\{0, 1\}$-rewards, each $r_i^{(t)} \in \{0, 1\}$ for $i \in \{1, 2\}$: a dominant expert gives a reward of 1.
- When there are only few changes of dominant experts in the two-expert case, the Be-The-Leader algorithm cannot change leaders promptly enough.
- A natural generalization [4] in which the regret is measured against a "dynamic" offline algorithm that can play different strategies in different rounds (under a small loss deviation).

Reward Maximizing Expert Setting: Modified Optimistic Multiplicative Updates with Predictable Sequences (2/2)

- Modified Optimistic Multiplicative Updates [4]
  1. Let $\bar{x}^{(1)} = x^{(1)} = \hat{x}^{(1)} = (\frac{1}{N}, \ldots, \frac{1}{N})^T$
  2. **for** $t \in [T]$ **do**

     Receive f_t and compute $\ell_t = \nabla f_t(\hat{x}^t)$.

     Update

     $\bar{x}_i^{(t+1)} = x_i^{(t)} e^{\eta r_i^{(t)}}$ with $\bar{Z}^{(t+1)} = \sum_j x_j^{(t)} e^{\eta r_j^{(t)}}$, and

     $x_i^{(t+1)} = (1-\beta)\bar{x}_i^{(t+1)} + \beta \frac{1}{N}$,

     $\hat{x}_i^{(t+1)} = x_i^{(t+1)} e^{\eta r_i^{(t)}}$ with $\bar{Z}^{(t+1)} = \sum_j x_j^{(t+1)} e^{\eta r_j^{(t)}}$.

     end for
- $\beta = \frac{1}{t}$
- Modified Double Multiplicative Updates: $M_{t+1} = \ell_t$

Few Changes of Dominant Experts

- k equal periods of time, each defined by its dominant expert: the dominant expert of a period is alternating between the two experts and $D_\infty = k$.
- At time step t, the ratio of the probability of the dominant expert to that of the other expert is at most ct for some constant c.
- Lemma: After the ith period of time steps, it takes at most $O(\log(T/k)/\eta)$ time steps for the ratio of the previous dominant expert's probability to the probability of the current dominant expert (of the $(i+1)$th period) to achieve at most 1.
- Theorem: The regret is at most $-T/2 + (k-1)\log(T/k)/\eta$. With $\eta = O(1/\sqrt{D_\infty}) = O(1/\sqrt{k})$, the regret is at most

$$-T/2 - \sqrt{k}\log(T/k) + K^{\frac{3}{2}}\log(T/k).$$

Conclusions and Future Work

- Optimistic (or double) lazy-update mirror descent
  - In each time step, a leader is strong compared with the other non-minimizers in terms of its much little current cumulative loss
  - The more frequent changes of leaders the more negative of the regret.
  - If the immediately previous loss vector has a small estimation error for the current loss vector, the regret can stay negative.
- Modified double-update multiplicative update algorithm of [4] for catching the switches of dominant experts quickly enough to obtain negative regrets.
- A "passive" market maker makes profit of a security by placing *limit orders* as *bids and asks* to reap the bid-ask spreads on the orderbook, yet takes exposure risk of the inventory.
  - A promising direction of market making design to try is to combine the merits of both the designs.

J. Abernethy, Y. Chen, and J. W. Vaughan.
Efficient market making via convex optimization, and a connection to online learning.
ACM Transactions on Economics and Computation, 1(1), 2012.

K. J. Arrow.
The role of securities in the optimal allocation of risk-bearing.
Review of Economic Studies, 31(2), 1964.

Y. Chen and J. W. Vaughan.
A new understanding of prediction markets via no-regret learning.
In *Proc. 11th ACM Conference on Electronic Commerce*, 2010.

C.-K. Chiang, C.-J. Lee, and C.-J. Lu.
Online learning in a gradually evolving world, 2011.
Unpublished manuscript.

C.-K. Chiang, T. Yang, C.-J. Lee, M. Mahdavi, C.-J. Lu, and S. Zhu.
Online optimization with gradual variations.

In *Proc. Conference on Learning Theory*, 2012.

A. Rakhlin and K. Sridharan.
Online learning with predictable sequences.
In *Proc. Conference on Learning Theory*, 2013.