CMSC 35401: The Interplay of Learning and Game Theory (Autumn 2022)

Prediction Markets (as a Forecasting Tool)

Instructor: Haifeng Xu

JOURNAL ARTICLE

Orange Juice and Weather

Richard Roll

The American Economic Review
Vol. 74, No. 5 (Dec., 1984),
pp. 861-880 (20 pages)

Futures of orange juice can be used to predict weather

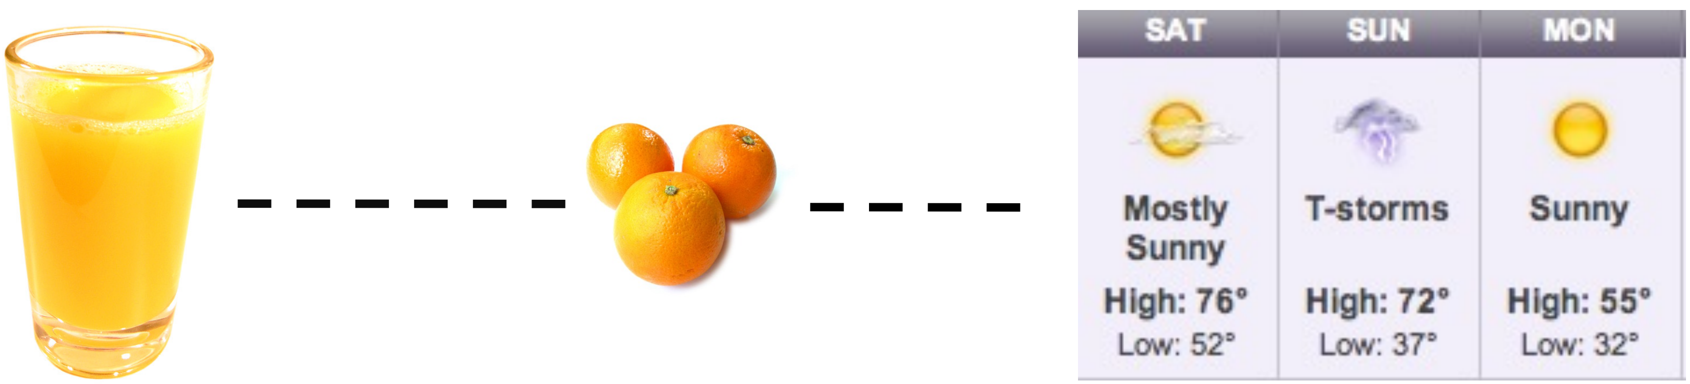

Outline

- Introduction to Prediction Markets
- Design of Prediction Markets
  - Logarithmic Market Scoring Rule (LMSR)
- LMSR and Exponential Weight Updates

Events of Interest for Prediction

- Will there be a HW3 for this course?
- Will UChicago faculties win a Nobel Prize in 2023?
- Will bit coin price exceed $25K tomorrow?
- Will Tesla's stock exceed $250 by the end of this year?
- Will the number of iPhones sold in 2022 exceed 200 million?
- Will there be a cure for cancer by 2025?
- Will the world be peaceful in 2050?
- . . .

The Prediction Problem

- An uncertain event to be predicted
  - E.g., will Tesla stock exceed $250 by Dec 2022?
- Dispersed information/evidence
  - E.g., Tesla employees, Tesla drivers, other EV company employees, government policy makers, etc.
- Goal: generate a prediction that is based on information aggregated from all sources
  - ML can also do prediction, but will see why markets have advantages

Bet ≈ Credible Opinion

Q: will P vs NP problem by solved by the end of 20'th century?

P vs NP would be solved by the end of the 20th century, if not sooner.
Betted one ounce of pure gold with

Michael Sipser

- Other examples: stock trading, gambling, . . .
- Betting intermediaries: Wall Street, Las Vegas, InTrade, . . .

Prediction Markets

A prediction market is a financial market that is designed for event prediction via information aggregation

- Payoffs of the traded financial contract are determined by outcomes of future events

$$\begin{cases} \$1 \text{ if UC wins Nobel} \\ \$0 \text{ otherwise} \end{cases}$$

A (financial) contract

Price of a contract? $1 × percentage of shares that bet on UC wining?

This is what we will be designing!

Prediction Markets: Examples

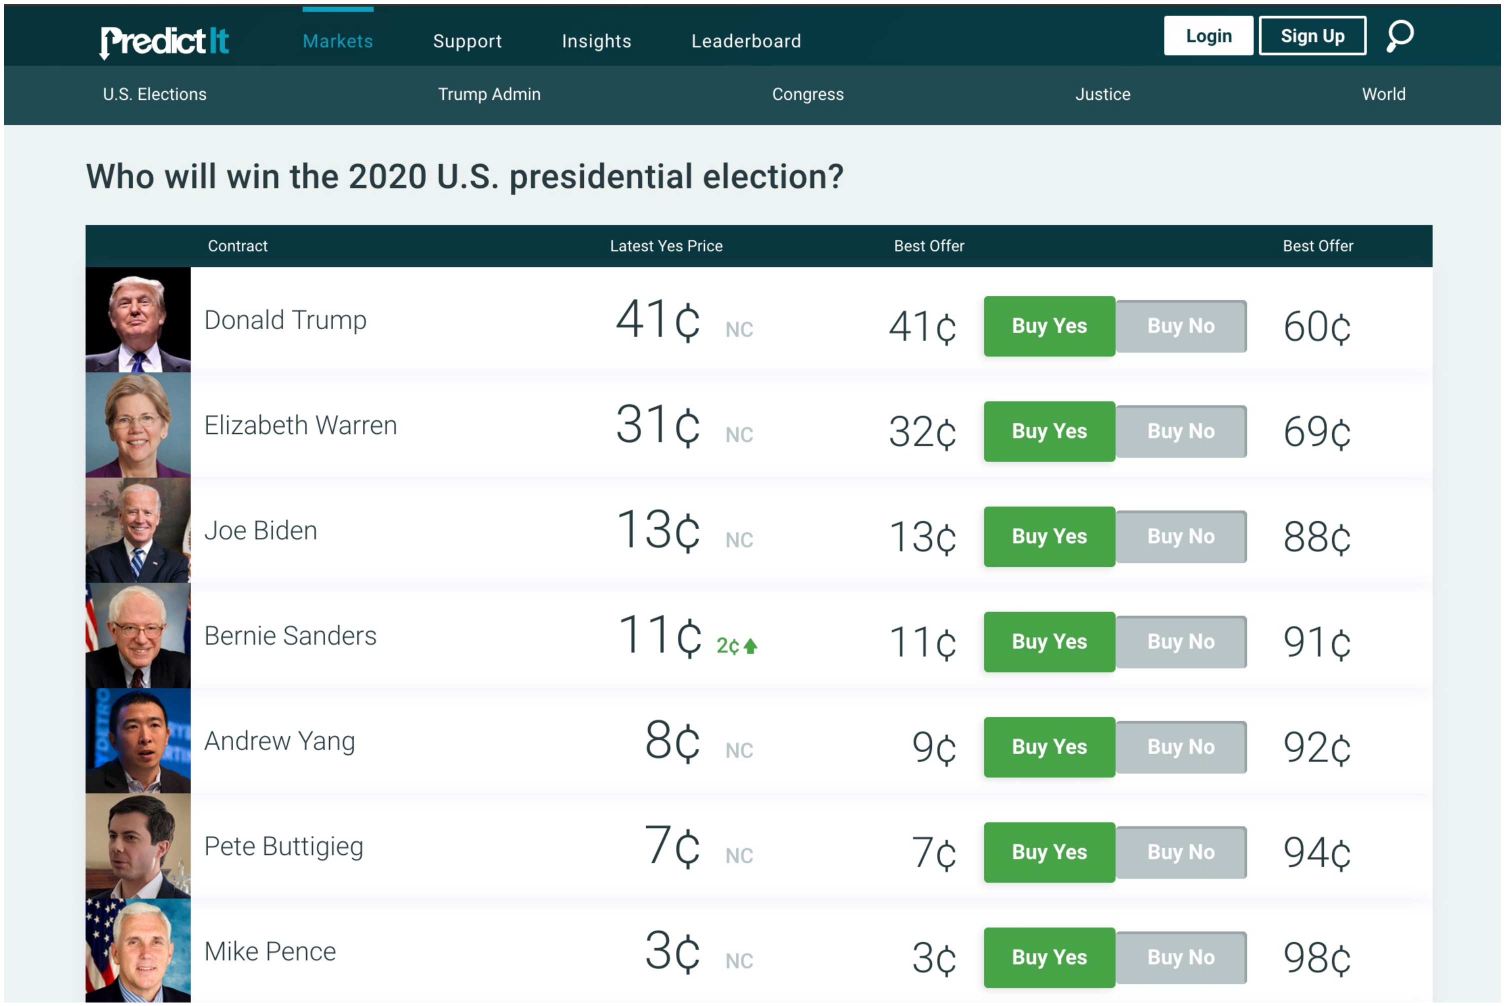

PredictIt   Markets   Support   Insights   Leaderboard   Login   Sign Up

U.S. Elections   Trump Admin   Congress   Justice   World

Who will win the 2020 U.S. presidential election?

Contract	Latest Yes Price		Best Offer			Best Offer
Donald Trump	41¢	NC	41¢	Buy Yes	Buy No	60¢
Elizabeth Warren	31¢	NC	32¢	Buy Yes	Buy No	69¢
Joe Biden	13¢	NC	13¢	Buy Yes	Buy No	88¢
Bernie Sanders	11¢	2¢	11¢	Buy Yes	Buy No	91¢
Andrew Yang	8¢	NC	9¢	Buy Yes	Buy No	92¢
Pete Buttigieg	7¢	NC	7¢	Buy Yes	Buy No	94¢
Mike Pence	3¢	NC	3¢	Buy Yes	Buy No	98¢

Prediction Markets: Examples

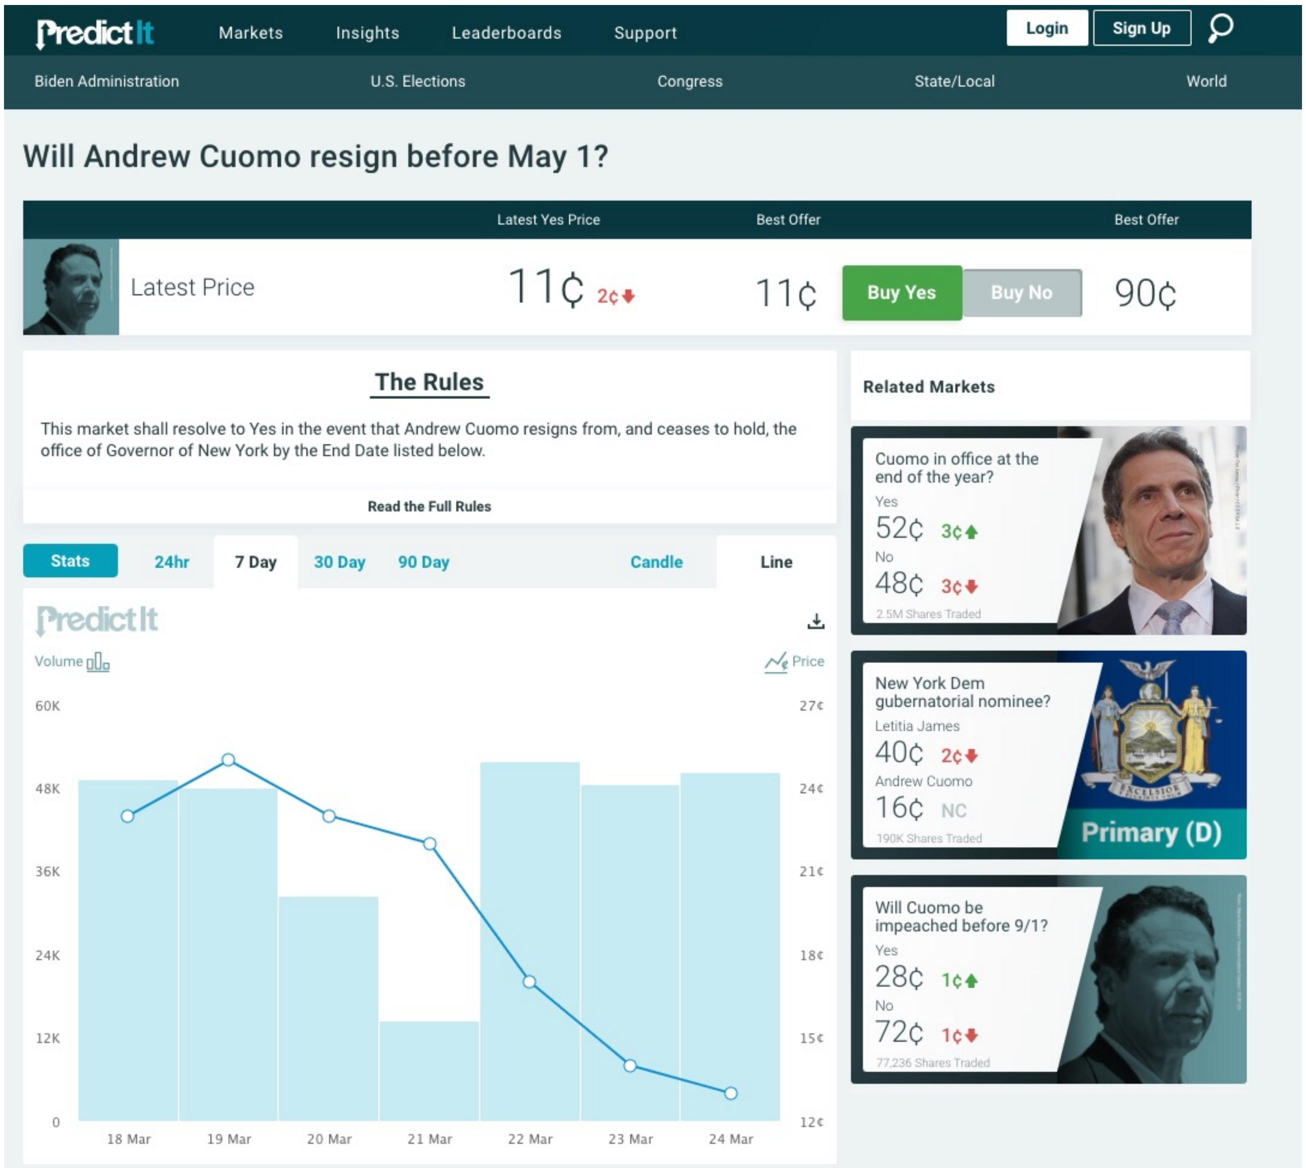

Prediction Markets: Examples

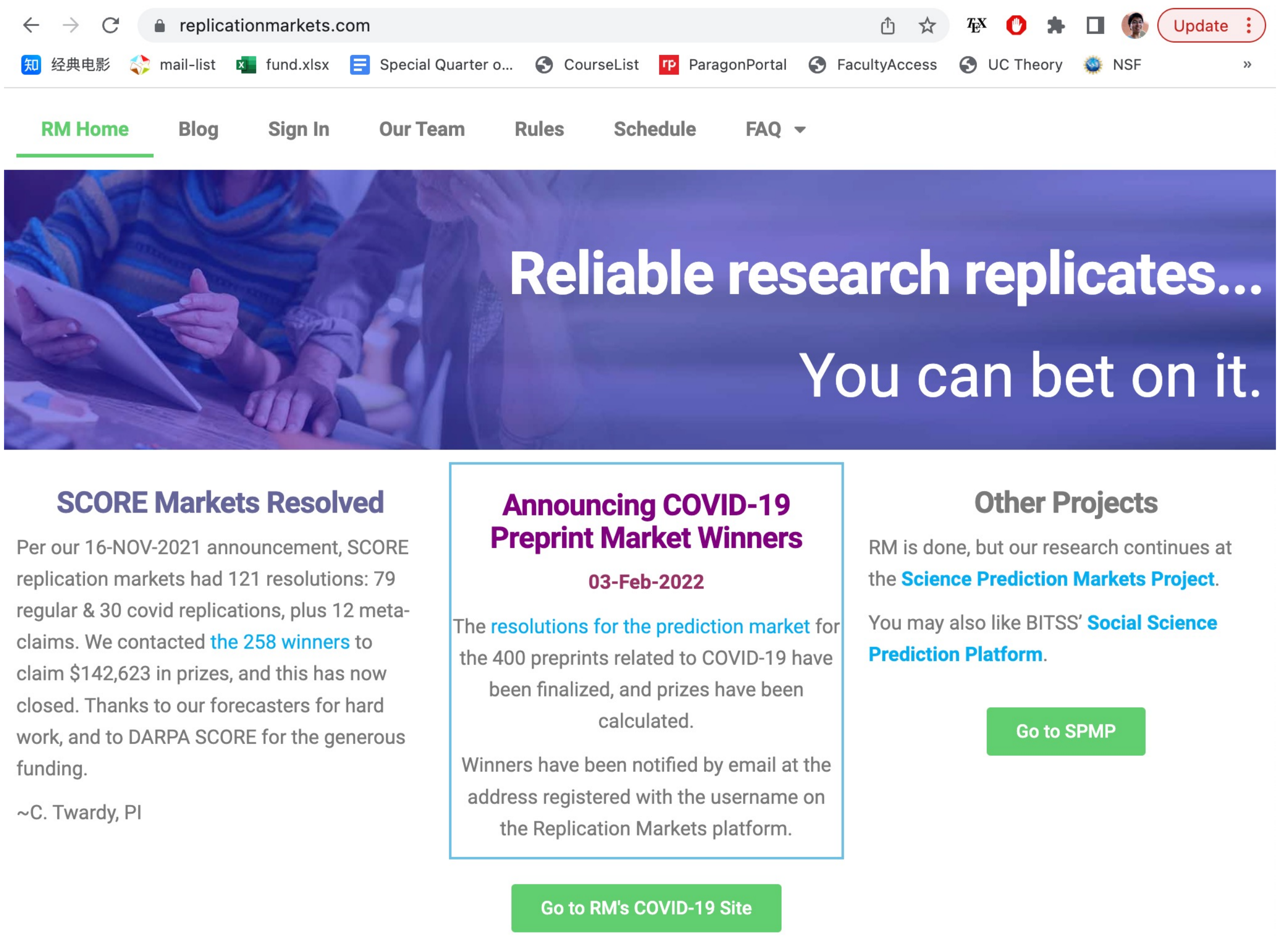

Replication Market

Prediction Markets: Examples

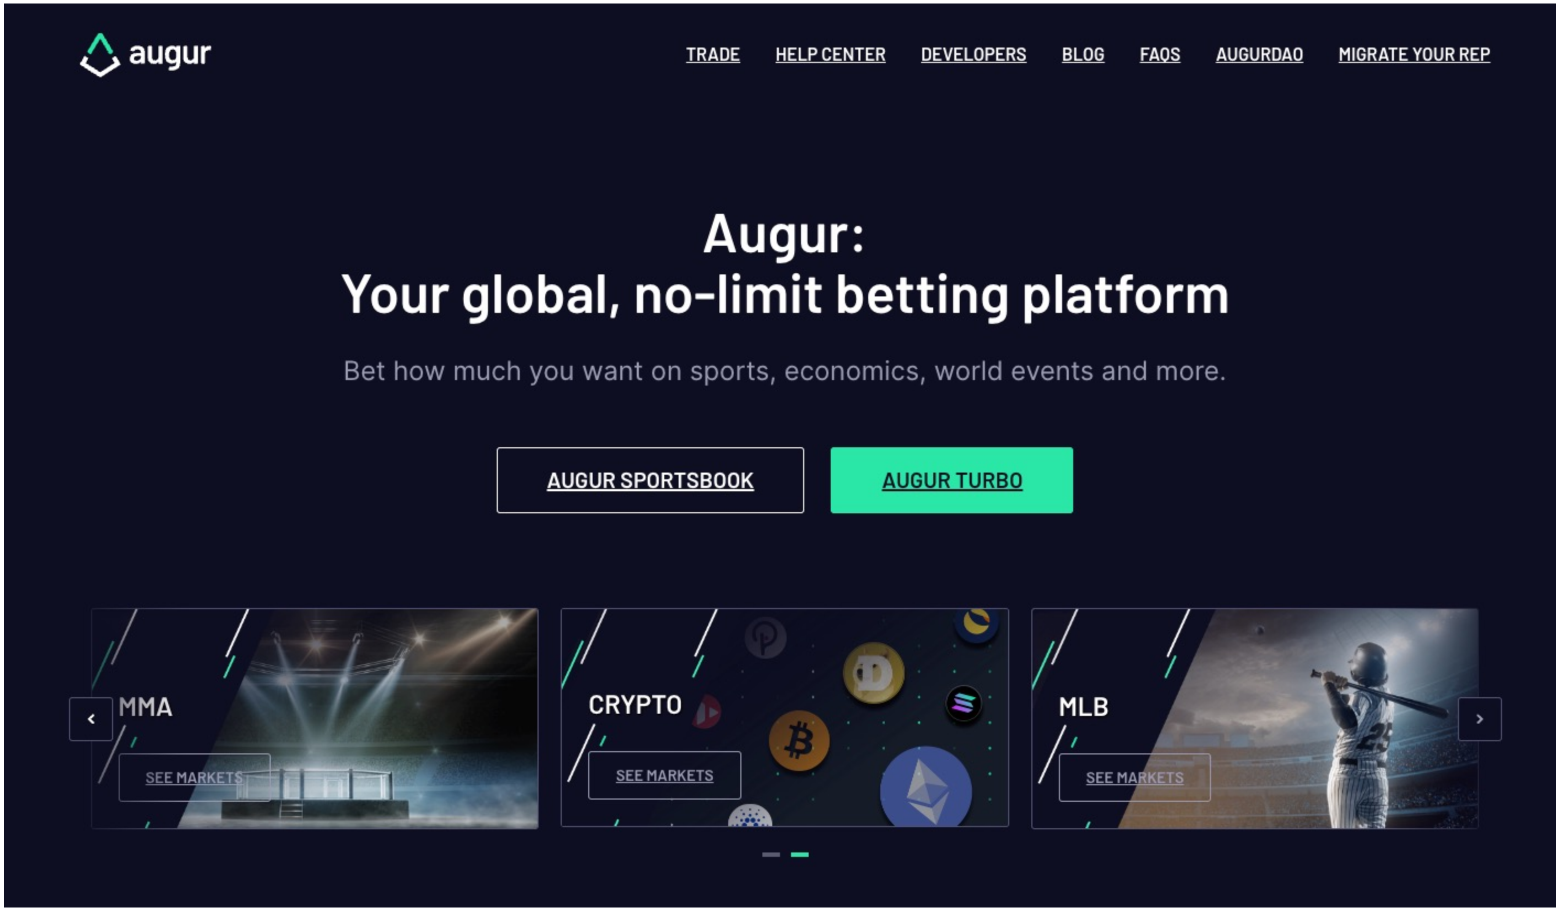

First decentralized prediction market (built on Ethereum blockchain)

Does It Work?

- Yes, evidence from real markets, lab experiments, and theory
  - I.E.M. beat political polls 451/596 [Forsythe 1992, 1999][Oliven 1995][Rietz 1998][Berg 2001][Pennock 2002]
  - HP market beats sales forecast 6/8 [Plott 2000]
  - Sports betting markets provide accurate forecasts of game outcomes [Gandar 1998][Thaler 1988][Debnath EC'03][Schmidt 2002]
  - Laboratory experiments confirm information aggregation [Plott 1982;1988;1997][Forsythe 1990][Chen, EC'01]
  - Theory: "rational expectations" [Grossman 1981][Lucas 1972]
  - More …

Why Can Markets Aggregate Information?

- Price $\approx Prob(\text{event}|\text{ all information})$

$1 if UC wins Nobel, $0 otherwise

Value of contract

?

Payoff	Event Outcome
$1	UC wins
$0	UC not

Why Can Markets Aggregate Information?

- Price $\approx Prob(\text{event}|\text{ all information})$

$1 if UC wins Nobel, $0 otherwise

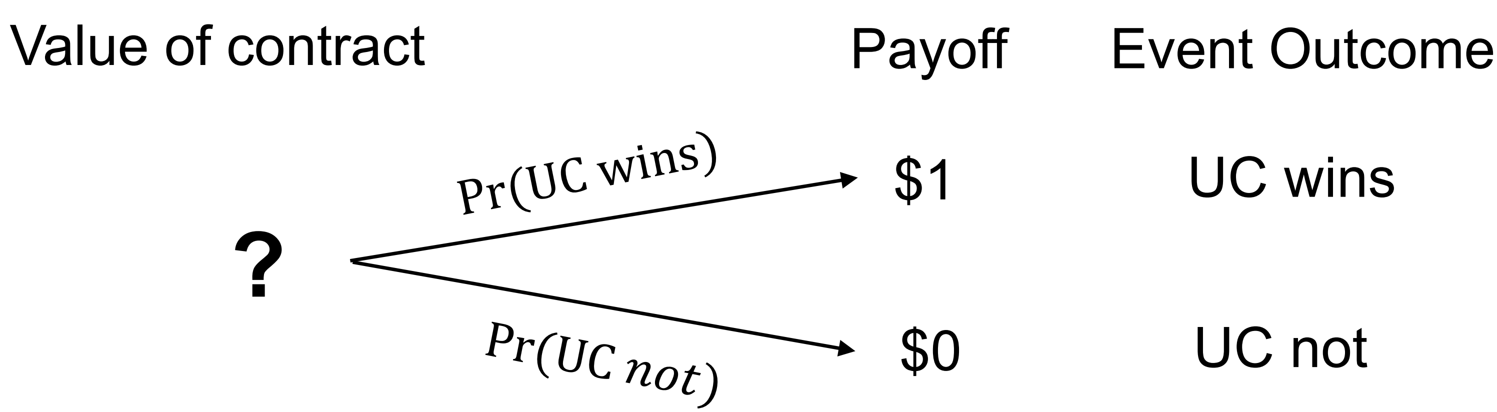

Why Can Markets Aggregate Information?

- Price ≈ $Prob(\text{event}|\text{ all information})$

$1 if UC wins Nobel, $0 otherwise

Value of contract		Payoff	Event Outcome
	$\Pr(\text{UC wins})$	$1	UC wins
$\Pr(\text{UC wins})$			
	$\Pr(\text{UC } not)$	$0	UC not

Value of contract ≈ P(UC wins)

Why Can Markets Aggregate Information?

- Price $\approx Prob(\text{event} |\ \text{all information})$

$1 if UC wins Nobel, $0 otherwise

Value of contract		Payoff	Event Outcome
	$\Pr(\text{UC wins})$	$1	UC wins
$\Pr(\text{UC wins})$			
	$\Pr(\text{UC } not)$	$0	UC not

Value of contract ≈ P(UC wins) ≈ Equilibrium price

Market Efficiency (a design goal)

Markets vs Other Prediction Approaches

Opinion Poll
- Sampling
- No incentive to be truthful
- Equally weighted information
- Hard to be real-time

Ask Experts
- Identifying experts can be hard
- Combining opinions is difficult

Prediction Markets
- Self-selection
- Monetary incentive and more
- Bet-weighted information
- Real-time
- Self-organizing

Other Prediction Approaches vs Markets

Machine Learning

- Historical data
- Assume past and future are related
- Hard to incorporate recent new information

Prediction Markets

- No need for data
- No assumption on past and future
- Immediately incorporate new information

Caveat: markets have their own problems too – manipulations, irrational traders, etc.

Outline

- Introduction to Prediction Markets
- Design of Prediction Markets (PMs)
    - Logarithmic Market Scoring Rule (LMSR)
- LMSR and Exponential Weight Updates

Some Design Objectives of PMs

Liquidity: people can find counterparties to trade whenever they want

Bounded loss: total loss of the market institution is bounded

Market efficiency: Price reflects predicted probabilities.

Computational efficiency: The process of operating the market should be computationally manageable.

Continuous Double Auction (CDA) Market

$1 if UC wins Nobel, $0 otherwise

- Buyer orders

$0.12

$0.09

$0.15

- Seller orders

$0.30

$0.17

$0.13

Continuous Double Auction (CDA) Market

$1 if UC wins Nobel, $0 otherwise

- Buyer orders
- Seller orders

Buyer orders	Seller orders
$0.15	**$0.13**
$0.12	**$0.17**
$0.09	**$0.30**

Continuous Double Auction (CDA) Market

$1 if UC wins Nobel, $0 otherwise

- Buyer orders
- Seller orders

$0.15

$0.12

$0.09

Price = $0.14

$0.13

$0.17

$0.30

What are Issues with CDA?

- Thin market problem
  - When there are not enough matches, trade may not happen.
- No trade theorem [Milgrom & Stokey 1982]
  - Why trade? These markets are zero-sum games (negative sum w/ transaction fees)
  - For all money earned, there is an equal (greater) amount lost; am I smarter than average?
  - Rational risk-neutral traders will *never* trade
  - But trade happens …

An Alternative: Market Maker (MM)

- A market maker is the market institution who sets the prices and is willing to accept orders (buy or sell) at the price specified.
- Why? Liquidity!
- Market makers bear risk. Thus, we want mechanisms that can bound the loss of market makers.

Example: Logarithmic Market Scoring Rule (LMSR [Hanson 03, 06])

- An automated market marker (MM)
- Sell or buy back contracts $1 iff e_1 · · · $1 iff e_n
- Described by a "value function" ($q = (q_1, \cdots, q_n) \in \mathbb{R}^n$ is sold #contracts)

$$V(q) = b \ln \sum_{j \in [n]} e^{q_j / b}$$

Parameter b adjusts liquidity

Example: Logarithmic Market Scoring Rule (LMSR [Hanson 03, 06])

- An automated market marker (MM)
- Sell or buy back contracts | \$1 iff e_1 | · · · | \$1 iff e_n |
- Described by a "value function" ($q = (q_1, \cdots, q_n) \in \mathbb{R}^n$ is sold #contracts)

$$V(q) = b \ln \sum_{j \in [n]} e^{q_j / b}$$

Parameter b adjusts liquidity

- To buy $x \in \mathbb{R}^n$ amount, a buyer "pays": $V(q+x) - V(q)$
  - Negative x_i's mean selling contracts to MM
  - Negative payment means market maker pays the buyer
  - Market starts with $V(0) = b \ln n$
- Price function

$$p_i(q) = \lim_{\delta \to 0} \frac{V(q + \delta e_i) - V(q)}{\delta}$$

Example: Logarithmic Market Scoring Rule (LMSR [Hanson 03, 06])

- An automated market marker (MM)
- Sell or buy back contracts

\$1 iff e_1	...	\$1 iff e_n

- Described by a "value function" ($q = (q_1, \cdots, q_n) \in \mathbb{R}^n$ is sold #contracts)

$$V(q) = b \ln \sum_{j \in [n]} e^{q_j / b}$$

Parameter b adjusts liquidity

- To buy $x \in \mathbb{R}^n$ amount, a buyer "pays": $V(q+x) - V(q)$
  - Negative x_i's mean selling contracts to MM
  - Negative payment means market maker pays the buyer
  - Market starts with $V(0) = b \ln n$
- Price function

$$p_i(q) = \lim_{\delta \to 0} \frac{V(q + \delta e_i) - V(q)}{\delta} = \frac{\partial V(q)}{\partial q_i} = \frac{e^{q_i / b}}{\sum_{j \in [n]} e^{q_j / b}}$$

Price Curve as a Function of Share Quantities

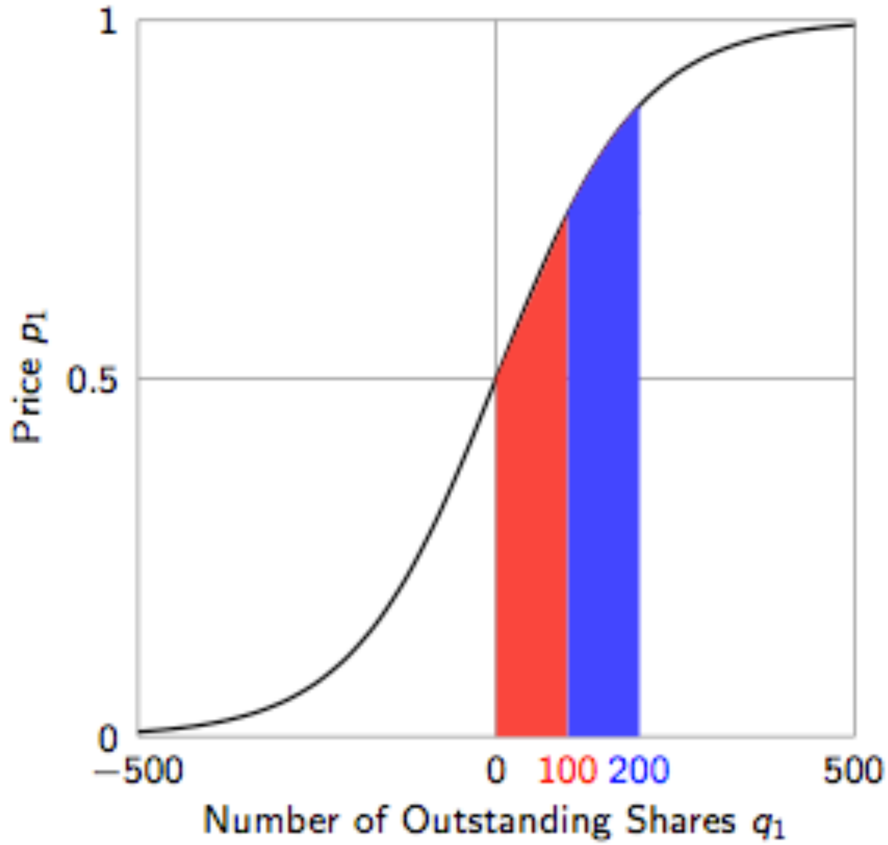

Example: Logarithmic Market Scoring Rule (LMSR [Hanson 03, 06])

- Value function $V(q) = b \ln \sum_{j \in [n]} e^{q_j / b}$

Q1: If your true belief of event $e_1, \cdots, e_n$ is $\lambda = (\lambda_1, \cdots, \lambda_n)$, how many shares of each contract should you buy?

- Say, you buy $x \in \mathbb{R}^n$ amount
- You pay $V(q + x) - V(q)$; Your expected return is $\sum_{j \in [n]} \lambda_j \cdot x_j$
- Expected utility is

$$U(x) = \sum_{j \in [n]} \lambda_j \cdot x_j - V(q + x) + V(q)$$

- Which x maximizes your utility?

$$0 = \frac{\partial U(x)}{\partial x_i} = \lambda_i - \frac{\partial V(q+x)}{\partial x_i}$$

Example: Logarithmic Market Scoring Rule (LMSR [Hanson 03, 06])

- Value function $V(q) = b \ln \sum_{j\in[n]} e^{q_j/b}$

Q1: If your true belief of event $e_1, \cdots, e_n$ is $\lambda = (\lambda_1, \cdots, \lambda_n)$, how many shares of each contract should you buy?

- Say, you buy $x \in \mathbb{R}^n$ amount
- You pay $V(q+x) - V(q)$; Your expected return is $\sum_{j\in[n]} \lambda_j \cdot x_j$
- Expected utility is

$$U(x) = \sum_{j\in[n]} \lambda_j \cdot x_j - V(q+x) + V(q)$$

- Which x maximizes your utility?

$$0 = \frac{\partial U(x)}{\partial x_i} = \lambda_i - \frac{\partial V(q+x)}{\partial x_i} = \lambda_i - \frac{e^{(q_i+x_i)/b}}{\sum_{j\in[n]} e^{(q_j+x_j)/b}}$$

This is the market price of contract i after your purchase

Example: Logarithmic Market Scoring Rule (LMSR [Hanson 03, 06])

- Value function $V(q) = b \ln \sum_{j \in [n]} e^{q_j/b}$

Q1: If your true belief of event $e_1, \cdots, e_n$ is $\lambda = (\lambda_1, \cdots, \lambda_n)$, how many shares of each contract should you buy?

Fact. The optimal amount you purchase is the amount that makes the market price equal to your belief λ. Your expected utility of purchasing this amount is always non-negative.

- Why non-negative?
  - Buy 0 amount leads to 0, so optimal amount is at least as good

Example: Logarithmic Market Scoring Rule (LMSR [Hanson 03, 06])

- Value function $V(q) = b \ln \sum_{j \in [n]} e^{q_j/b}$

Q1: If your true belief of event $e_1, \cdots, e_n$ is $\lambda = (\lambda_1, \cdots, \lambda_n)$, how many shares of each contract should you buy?

Fact. The optimal amount you purchase is the amount that makes the market price equal to your belief λ. Your expected utility of purchasing this amount is always non-negative.

- This is the expected utility you believe, but may be incorrect since your λ may be inaccurate!
  - So, buy only when your prediction is really more accurate than the current market prediction
  - Achieves market efficiency: price = current best market prediction

Example: Logarithmic Market Scoring Rule (LMSR [Hanson 03, 06])

- Value function $V(q) = b \ln \sum_{j\in[n]} e^{q_j/b}$

Q2: If market ends up with $q = (q_1, \cdots, q_n)$ shares for each contract, how much money did the MM collect?

- Answer: $V(q) - V(0) = V(q) - b \ln n$
- But after event outcome is realized, MM need to pay based on contracts – what is the worst-case loss of MM?

Example: Logarithmic Market Scoring Rule (LMSR [Hanson 03, 06])

- Value function $V(q) = b \ln \sum_{j \in [n]} e^{q_j/b}$

Fact. After event outcome realizes and MM pays the contract, worst case MM loses is $b \ln n$ (i.e., bounded).

Proof

- Only one event will be realized, say it is event e_i
- MM utility is $V(q) - b \ln\, n - q_i$

$$\geq b \ln e^{q_i/b} - b \ln n - q_i$$

$$\geq q_i - b \ln\, n - q_i$$

$$\geq -b \ln\, n$$

"=" can be achieved by letting $q_i \to \infty$ (i.e., people all think e_i will occur)

Example: Logarithmic Market Scoring Rule (LMSR [Hanson 03, 06])

- Has been implemented by several prediction markets
  - E.g., InklingMarkets, Washington Stock Exchange, BizPredict, Net Exchange, and (reportedly) at YooNew.

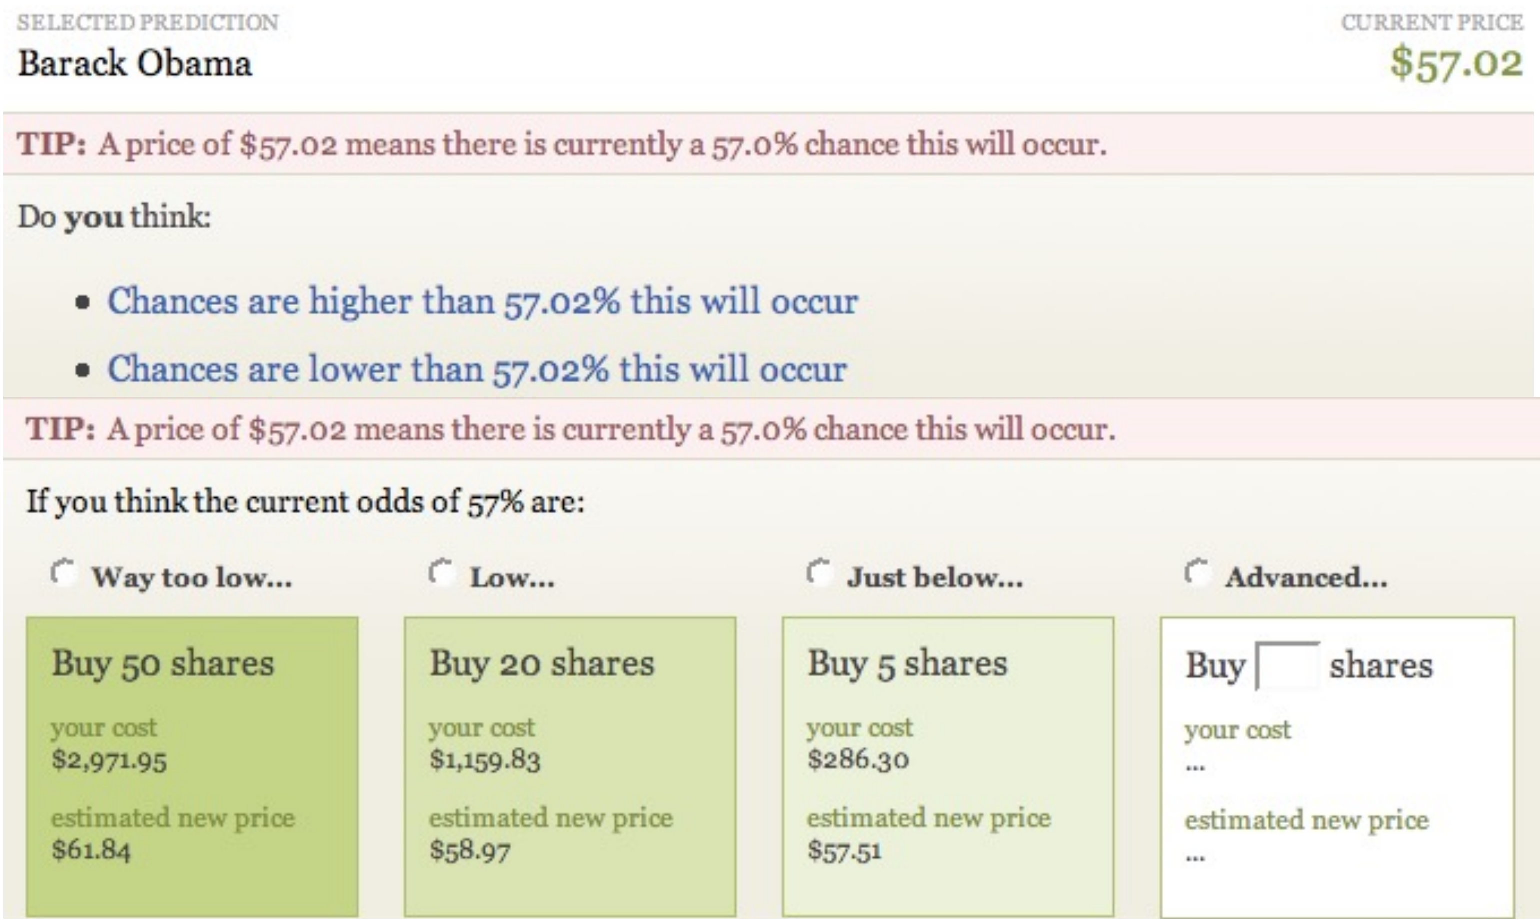

Outline

- Introduction to Prediction Markets
- Design of Prediction Markets
  - Logarithmic Market Scoring Rule (LMSR)
- LMSR and Exponential Weight Updates (EWU)

Recap: Exponential Weight Update

- Played for T rounds; each round selects an action $i \in [n]$
- Maintains weights over n actions: $w_t(1), \cdots, w_t(n)$
- Observe cost vector c_t, and update $w_{t+1}(i) = w_t(i) \cdot e^{-\epsilon c_t(i)}, \forall i \in [n]$

Action 1, $w_t(1)$

Action 2, $w_t(2)$

. . .

Action n, $w_t(n)$

$$w_{t+1}(i) = w_t(i) \cdot e^{-\epsilon c_t(i)}$$

$$= [w_{t-1}(i) \cdot e^{-\epsilon c_{t-1}(i)}] \cdot e^{-\epsilon c_t(i)}$$

$$= \cdots = e^{-\epsilon C_t(i)} \text{ where } C_t(i) = \sum_{\tau \le t} c_\tau(i)$$

Recap: Exponential Weight Update

- Played for T rounds; each round selects an action $i \in [n]$
- Maintains weights over n actions: $w_t(1), \cdots, w_t(n)$
- Observe cost vector c_t, and update $w_{t+1}(i) = w_t(i) \cdot e^{-\epsilon c_t(i)}, \forall i \in [n]$
- At round $t + 1$, select action i with probability

$$\frac{w_t(i)}{W_t} = \frac{e^{-\epsilon C_t(i)}}{\sum_{j\in[n]} e^{-\epsilon C_t(j)}}$$

where $C_t = \sum_{\tau \le t} c_t$ is the accumulated cost vector

This looks very much like the price function in LMSR (q is the accumulated sales quantity)

$$p_i = \frac{e^{q_i/b}}{\sum_{j\in[n]} e^{q_j/b}}$$

EWU vs LMSR

- Exponential Weight Update
  - n actions
  - Maintain weight $w_t(i)$
  - Total cost $C_T(i) = \sum_{t \le T} c_t(i)$
  - Select i with prob

$$p_i = \frac{e^{-\epsilon C_t(i)}}{\sum_{j \in [n]} e^{-\epsilon C_t(j)}}$$

  - Weights reflect how good an action is
  - Care about worst case regret

$$C_T(\text{Alg}) - \min_i C_T(i)$$

- LMSR
  - n contracts (i.e., outcomes)
  - Maintain prices $p(i)$
  - Total shares sold $q(i)$
  - Price of contract i

$$p_i = \frac{e^{q_i/b}}{\sum_{j \in [n]} e^{q_j/b}}$$

  - Prices reflect how probable is an event
  - Care about worst case MM loss

$$(\text{\$ received}) - \min_i q(i)$$

Remarks

- LMSR is just one particular automatic MM
- Similar relation holds for other market markers and no-regret learning algorithms (see [Chen and Vaughan 2010])
- Markets can potentially be a very effective forecasting tool
  - Big project: "replication market" for DARPA SCORE program

DARPA DEFENSE ADVANCED RESEARCH PROJECTS AGENCY

EXPLORE BY TAG

ABOUT US / OUR RESEARCH / NEWS / EVENTS / WORK WITH US /

Defense Advanced Research Projects Agency > Program Information

Systematizing Confidence in Open Research and Evidence (SCORE)

Dr. Adam Russell

The Department of Defense (DoD) often leverages social and behavioral science (SBS) research to design plans, guide investments, assess outcomes, and build models of human social systems and behaviors as they relate to national security challenges in the human domain. However, a number of recent empirical studies and meta-analyses have revealed that many SBS results vary dramatically in terms of their ability to be independently reproduced or replicated, which could have real-world implications for DoD's plans, decisions, and models. To help address this situation, DARPA's Systematizing Confidence in Open Research and Evidence (SCORE) program aims to develop and deploy automated tools to assign "confidence scores" to different SBS research results and claims. Confidence scores are quantitative measures that should enable a DoD consumer of SBS research to understand the degree to which a particular claim or result is likely to be reproducible or replicable. These tools will assign explainable confidence scores with a reliability that is equal to, or better than, the best current human expert methods. If successful, SCORE will enable DoD personnel to quickly calibrate the level of confidence they should have in the reproducibility and replicability of a given SBS result or claim, and thereby

Remarks

- LMSR is just one particular automatic MM
- Similar relation holds for other market markers and no-regret learning algorithms (see [Chen and Vaughan 2010])
- Markets can potentially be a very effective forecasting tool
  - Big project: “replication market” for DARPA SCORE program
- Mechanism design for prediction tasks
  - ML is not the only way of making predictions
  - But markets and ML may augment each other…

Thank You

Haifeng Xu

University of Chicago

haifengxu@uchicago.edu