CS 598:
AI Methods for Market Design

Lecture 10: Prediction Markets

Xintong Wang

Spring 2024

Recap

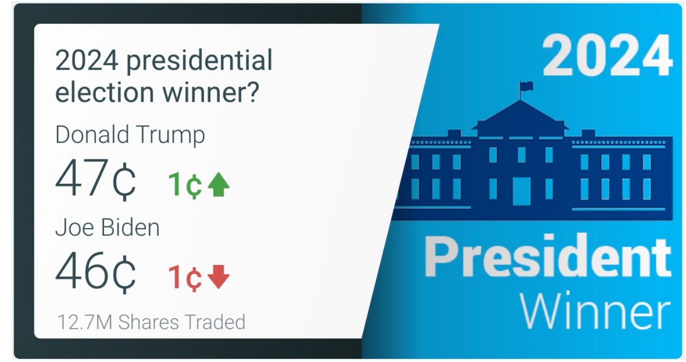

- Eliciting beliefs about something verifiable in the future
  - E.g., Will Trump be the 2024 presidential election winner?
  - Scoring rules & prediction markets
- Eliciting information without (easy) verification
  - E.g., Does a plumber do high quality work?
    Is a restaurant good for friend gatherings?
  - Scoring rules & peer prediction

How to crowdsource information to make reliable predictions?

Recap: Scoring Rules
(How to Pay a Forecaster)

- Possible outcomes $O = \{o_0, \dots, o_{m-1}\}$, indexed by k
- An agent's *true belief* p
  - E.g., I believe it will rain tomorrow with probability 0.5
- An agent's *belief report* q
- A **scoring rule** pays $s(q, o_k)$ if the outcome is o_k
  - The payment is contingent on the outcome
- Expected payment

$$E_{o\sim p}[s(q, o)] = \sum_k p_k \cdot s(q, o_k)$$

Recap: Scoring Rules (How to Pay a Forecaster)

- A scoring rule is strictly proper if, for every belief p, the expected payment

$$E_{o\sim p}[s(q,o)] = \sum_k p_k \cdot s(q,o_k)$$

is *uniquely maximized* through truthful report (q=p)

Example 1: Linear Scoring Rule

- The weather for tomorrow is a random variable W
- The outcome space is {sun, rain}
- True belief p = Pr(W=rain)
- Reported belief q
- Linear scoring rule: $s_{linear}(q, o_k) = q_k$
  - If it rains, then pay q; if it is sunny, then pay 1-q
- What is the expected payment?

p*q + (1-p)*(1-q)

  - Suppose p=0.6. What is the best report? q=1
  - Based on p, an agent will only report $q \in \{0, 1\}$

Example 2: Logarithmic Scoring Rule

- Logarithmic scoring rule

$$s_{log}(q, o_k) = \ln(q_k)$$

- Expected payment under weather forecasting

p*ln(q)+ (1-p)*ln(1-q)

- Verify optimality
  - First-order: p/q+1/(q-1)-p/(q-1) = 0 → q=p
  - Second-order derivative is negative
- Logarithmic scoring rule is strictly proper

Prediction Market

- A market designed for information aggregation
- Agents can "bet on beliefs", by trading contracts whose payoffs associated with an observed outcome in the future

Prediction Market as a Forecasting Tool

Goal: Produce a forecast based on information dispersed among agents from all sources

[Hypermind, December 2020]

Prediction Market as a Forecasting Tool

- Construct a contract on an outcome (e.g., time of approval)

Q1, 2021 (or before)	1¢
Q2, 2021	27¢
Q3, 2021	55¢
Q4, 2021 (or later)	17¢

Aggregate information via agents trading *contracts* written on event outcomes

Payoff = { $1 if FDA approves one in Q2, 2021; $0 otherwise

[Hypermind, December 2020]

Prediction Market as a Forecasting Tool

- Bet on beliefs (buy if price < \$p, and sell if price >\$p)

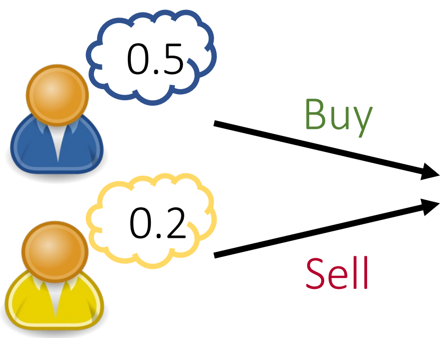

Q1, 2021 (or before)	1¢
Q2, 2021	27¢
Q3, 2021	55¢
Q4, 2021 (or later)	17¢

Aggregate information via agents trading *contracts* written on event outcomes

Payoff = { \$1 if FDA approves one in Q2, 2021; \$0 otherwise }

[Hypermind, December 2020]

Prediction Market as a Forecasting Tool

- Price represents aggregated belief, given dispersed information

Q1, 2021 (or before)	1¢
Q2, 2021	27¢
Q3, 2021	55¢
Q4, 2021 (or later)	17¢

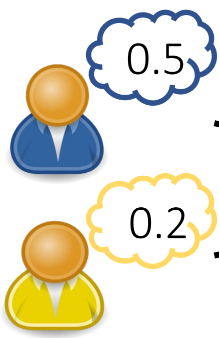

Buy

Sell

Aggregate information via agents trading *contracts* written on event outcomes

Payoff = { $1 if FDA approves one in Q2, 2021; $0 otherwise

Price of contract ≈ Prob (event | all information)

[Hypermind, December 2020]

Prediction Landscape

	1 person	n >1 people
(verification) elicit belief	scoring rule	prediction market
(no verification) elicit signal	x	peer prediction

Other Prediction Methods vs. Prediction Market

Opinion Poll

- Sample with equally weighted inputs
- No incentive to be truthful
- Hard to be real-time

Ask Experts

- Hard to identify experts
- Hard to combine information

Machine Learning

- Need historical data, assuming past and future are related
- Hard to incorporate new information

Prediction Market

- Self-selection with bet-weighted inputs
- Monetary incentive
- No need for (assumptions on) data
- Real-time with new information immediately incorporated

Financial vs. Prediction Market

	Financial	**Prediction**
Primary Use	Capital allocation Hedge risk	Information aggregation
Secondary Use	Information aggregation	Hedge risk

Applications

- PredictIt, Iowa Electronic Markets
- Google, Ford, HP, etc.: user internal prediction markets for sales forecasts (software by firms, e.g., CultivateLabs)
- CMU Gates-Hillman prediction market
- Hollywood Stock Exchange (HSX)
- Prosper: blockchain-based prediction markets

Market Designs

- Design 1: continuous double auction (CDA)

PredictIt, Iowa Electronic Market, HSX

- Design 2: automated market maker (AMM) using market scoring rule

CultivateLabs, Prosper (Ethereum smart contract), DeFi such as Uniswap

Some Desirable Properties

- Liquidity (can always trade *any* quantity)
- Information aggregation
- Real-time
- No "round-trip" *arbitrage* (profit at no risk)
- Bounded loss for the market designer

Continuous Double Auction (CDA)

- Limit order book

	Price	Shares
Sell Orders	…	…
	100.12	4
	100.10	15
	100.04	20
	100.03	8
Buy Orders	100.01	3
	99.99	11
	99.98	18
	99.95	20
	…	…

Price	Shares
…	…
100.12	4
100.10	15
100.04	20
100.03	8
100.01	3
100.00	2
99.99	11
99.98	18
…	…

Submit
"Buy 2 shares @ $100.00"

Price	Shares
…	…
100.12	4
100.10	15
100.04	20
100.03	8
100.00	2
99.99	11
99.98	18
99.95	20
…	…

Submit
"Sell 3 shares @ $100.01"

Price	Shares
…	…
100.12	4
100.10	15
100.04	20
100.03	8
99.99	11
99.98	18
99.95	20
99.91	34
…	…

Cancel
"Buy 2 shares @ $100.00"

Continuous Double Auction (CDA)

- CDAs are real-time, but can have low liquidity

Will 2019 be a warmer year than 2009?

Contract	Bid	Ask	Last	Vol	Chge
GLOBALTEMP.2019>2009 Global Average Temperature for 2019 to be higher than for 2009 M Trade	30.0	-	95.0	1	0

Will 2019 be 0.2 degrees celsius warmer than 2009?

Contract	Bid	Ask	Last	Vol	Chge
GLOBALTEMP.2019.0.2C>2009 Global Average Temperature for 2019 to be 0.2 degrees celsius (or more) higher than for 2009 M Trade	30.0	96.0	98.0	1	0

(Das)

Call Market

- Buy orders (over T)

0.15

0.12

0.09

0.05

- Sell orders (over T)

0.08

0.11

0.13

0.17

0.30

Orders are batched together and matched at predetermined time intervals

Somewhat solve thin market problem, but not real-time

Call Market

• Buy orders (over T)	• Sell orders (over T)	
0.15	0.08	*Two trades with price in [0.11, 0.12]*
0.12	0.11	
0.09	0.13	
0.05	0.17	
	0.30	

Orders are batched together and matched at predetermined time intervals

Somewhat solve thin market problem, but not real-time

Automated Market Maker (AMM)

- Quote prices and offer to trade any quantity
- Goal: improve liquidity, and thus information aggregation

Automated Market Maker (AMM)

- Quote prices and offer to trade any quantity
- Goal: improve liquidity, and thus information aggregation
- Will Rutgers appear in NCAA tournament 2025?

Market State: x

	Yes	No
	0	0
Buy 2 for Yes	2	0
Buy 5 for Yes	7	0
Buy 2 for No	7	2
Sell 1 for Yes	6	2

How to charge these trades?

Some Desirable Properties (AMM)

- No “round-trip” arbitrage
- Prices nonnegative, sum to one (i.e., =probability)
- Responsiveness (i.e., if buy then price increases; if sell then price decreases)
- Liquidity (i.e., relatively small price change after a small trade)
- Myopic incentives (i.e., trade until price=belief)
- Bounded loss to the market maker

Cost-Function-Based AMM

- Cost function (convex, strictly increasing): $C(x)$

  Example: $C(x) = \beta \ln \left(\sum_{j=0}^{m-1} e^{x_j/\beta} \right)$

- Will Rutgers appear in NCAA tournament 2025?

  Example: $C(x_0, x_1) = \beta \ln(e^{\frac{x_0}{\beta}} + e^{\frac{x_1}{\beta}})$

  Market State: x

	Yes	No	
	0	0	
Buy 2 for Yes	2	0	Trader pays C(2, 0) - C(0, 0)
Buy 5 for Yes	7	0	Trader pays C(7, 0) - C(2, 0)
Buy 2 for No	7	2	Trader pays C(7, 2) - C(7, 0)
Sell 1 for Yes	6	2	Trader pays C(6, 2) - C(7, 2) ← Negative

Cost-Function-Based AMM

- Cost function (convex, strictly increasing): $C(x)$

Example: $C(x) = \beta \ln \left(\sum_{j=0}^{m-1} e^{x_j/\beta} \right)$

- Will Rutgers appear in NCAA tournament 2025?

Example: $C(x_0, x_1) = \beta \ln(e^{\frac{x_0}{\beta}} + e^{\frac{x_1}{\beta}})$

Market State: x

	Yes	No
	0	0
Buy 2 for Yes	2	0
Buy 5 for Yes	7	0
Buy 2 for No	7	2
Sell 1 for Yes	6	2

No round-trip arbitrage!

AMM gets $C(x^{(T)}) - C(x^{(0)})$ & pays $1 to winners!

Trader pays C(2, 0) - C(0, 0)

Trader pays C(7, 0) - C(2, 0)

Trader pays C(7, 2) - C(7, 0)

Trader pays C(6, 2) - C(7, 2) ← Negative

Cost-Function-Based AMM

- Analyze the cost function: $C(x_0, x_1) = \beta \ln(e^{\frac{x_0}{\beta}} + e^{\frac{x_1}{\beta}})$
- Price for an infinitesimal amount: $\pi_k(x) = \frac{\partial}{\partial x_k} C(x)$
  - Price for "YES": $\pi_0(x) = \frac{e^{x_0/\beta}}{e^{x_0/\beta} + e^{x_1/\beta}}$
  - Price for "NO": $\pi_1(x) = \frac{e^{x_1/\beta}}{e^{x_0/\beta} + e^{x_1/\beta}}$

- *Does this look familiar?*

Some Desirable Properties (AMM)

- ✓ No "round-trip" arbitrage
- ✓ Prices nonnegative, sum to one (i.e., =probability)
- ✓ Responsiveness (i.e., if buy then price increases; if sell then price decreases)
- ✓ Liquidity (i.e., relatively small price change after a small trade) *More liquid as beta is larger*
- Myopic incentives (i.e., trade until price=belief)
- Bounded loss to the market maker

Cost-Function-Based AMM

- Myopic incentives: optimal for an agent to trade until instantaneous price $\pi = p$ (agent belief)

- Connect to sequential logarithmic scoring rule
  - Initialize the market: $q^{(0)}$ is uniform
  - Sequence of reports: $q^{(0)}, q^{(1)}, \ldots, q^{(n)}$
  - Upon realization of o_k, the *ith agent* pays

$$s\left(q^{(i-1)}, o_k\right) - s\left(q^{(i)}, o_k\right)$$

  - Take *s* to be log scoring rule, i.e., $s_{log}(q, o_k) = \beta \ln(q_k)$. Is it rational to report truthfully in position *i*?

Cost-Function-Based AMM

- Myopic incentives: optimal for an agent to trade until instantaneous price $\pi = p$ (agent belief)

- Connect to sequential logarithmic scoring rule
  - Initialize the market: $q^{(0)}$ is uniform
  - Sequence of reports: $q^{(0)}, q^{(1)}, \ldots, q^{(n)}$
  - Upon realization of o_k, the *ith agent* pays
    $$s\left(q^{(i-1)}, o_k\right) - s(q^{(i)}, o_k)$$
  - Take *s* to be log scoring rule, i.e., $s_{log}(q, o_k) = \beta \ln(q_k)$. Is it rational to report truthfully in position *i*? YES!

Cost-Function-Based AMM

- Myopic incentives: optimal for an agent to trade until instantaneous price $\pi = p$ (agent belief)

- Connect to sequential logarithmic scoring rule
  - Initialize the market: $q^{(0)}$ is uniform
  - Sequence of reports: $q^{(0)}, q^{(1)}, \dots, q^{(n)}$
  - Upon realization of o_k, the *ith agent* pays
    $$s\left(q^{(i-1)}, o_k\right) - s(q^{(i)}, o_k)$$
  - Take *s* to be log scoring rule, i.e., $s_{log}(q, o_k) = \beta \ln(q_k)$. Is it rational to report truthfully in position *i*? YES!
  - The worst-case total cost = $s\left(q^{(n)}, o_k\right) - s(q^{(0)}, o_k)$
    $$\leq \beta \ln(1) - \beta \ln(1/m) = \beta \ln(m)$$

Some Desirable Properties (AMM)

- ✓ No "round-trip" arbitrage
- ✓ Prices nonnegative, sum to one (i.e., =probability)
- ✓ Responsiveness (i.e., if buy then price increases; if sell then price decreases)
- ✓ Liquidity (i.e., relatively small price change after a small trade) *More liquid as beta is larger*
- ✓ Myopic incentives (i.e., trade until price=belief)
- ✓ Bounded loss to the market maker

Cost-Function-Based AMM

- Will Rutgers appear in NCAA tournament 2025?

$\beta = 1, \mathrm{C}(x) = \ln(e^{x_0} + e^{x_1}), s_{log}(q, o_k) = \ln(q_k)$

	Yes	No	Payment	$\pi(Yes)$	$\pi(No)$	Payment | Yes	Payment | No
	0	0	–	0.5	0.5	–	
Buy 1 for Yes	1	0	0.62 $\ln(e^1 + e^0) - \ln(e^0 + e^0)$	0.73 $e^1/(e^1 + e^0)$	0.27	*-0.38* ln(0.5)-ln(0.73)	0.62 ln(0.5)-ln(0.27)
Buy 2 for Yes	3	0	1.73 $\ln(e^3 + e^0) - \ln(e^1 + e^0)$	0.95 $e^3/(e^3 + e^0)$	0.05	-0.26 ln(0.73)-ln(0.95)	1.73 ln(0.27)-ln(0.05)
Buy 1 for No	3	1	0.08 $\ln(e^3 + e^1) - \ln(e^3 + e^0)$	0.88 $e^3/(e^3 + e^1)$	0.12	0.08 ln(0.95)-ln(0.88)	-0.92 ln(0.05)-ln(0.12)

Summary: Scoring-Rule based AMM

- Cost-function-based AMM, with cost function

$$C(x) = \beta \ln \left(\sum_{j=0}^{m-1} e^{x_j/\beta} \right)$$

Logarithmic market scoring rule (LMSR) AMM

- Satisfy all desirable properties!
- Used by CultivateLabs, Prosper, …

How about these scenarios?

- Payoff is function of common variables, e.g. 50 states elect Dem or Rep

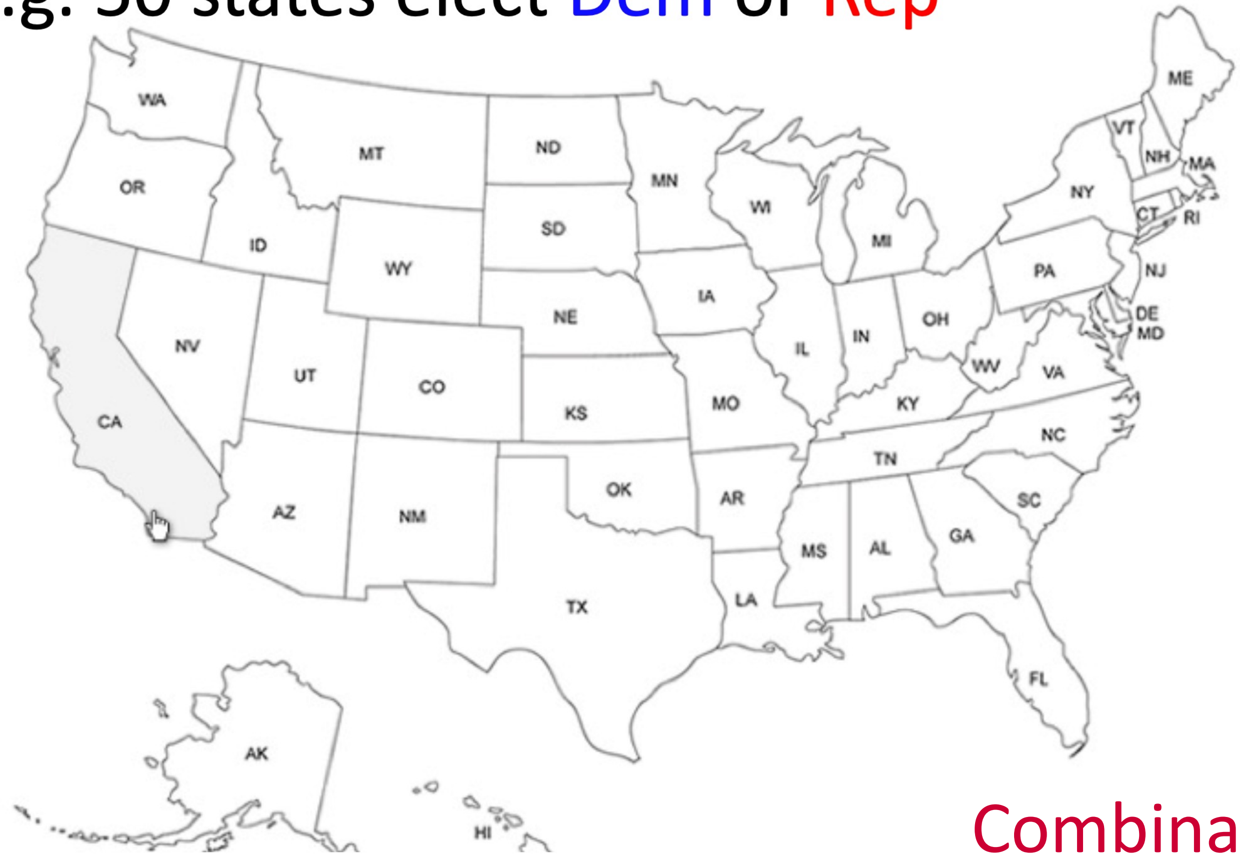

Combinatorial

How about these scenarios?

Large outcome space / even continuous

How about these scenarios?

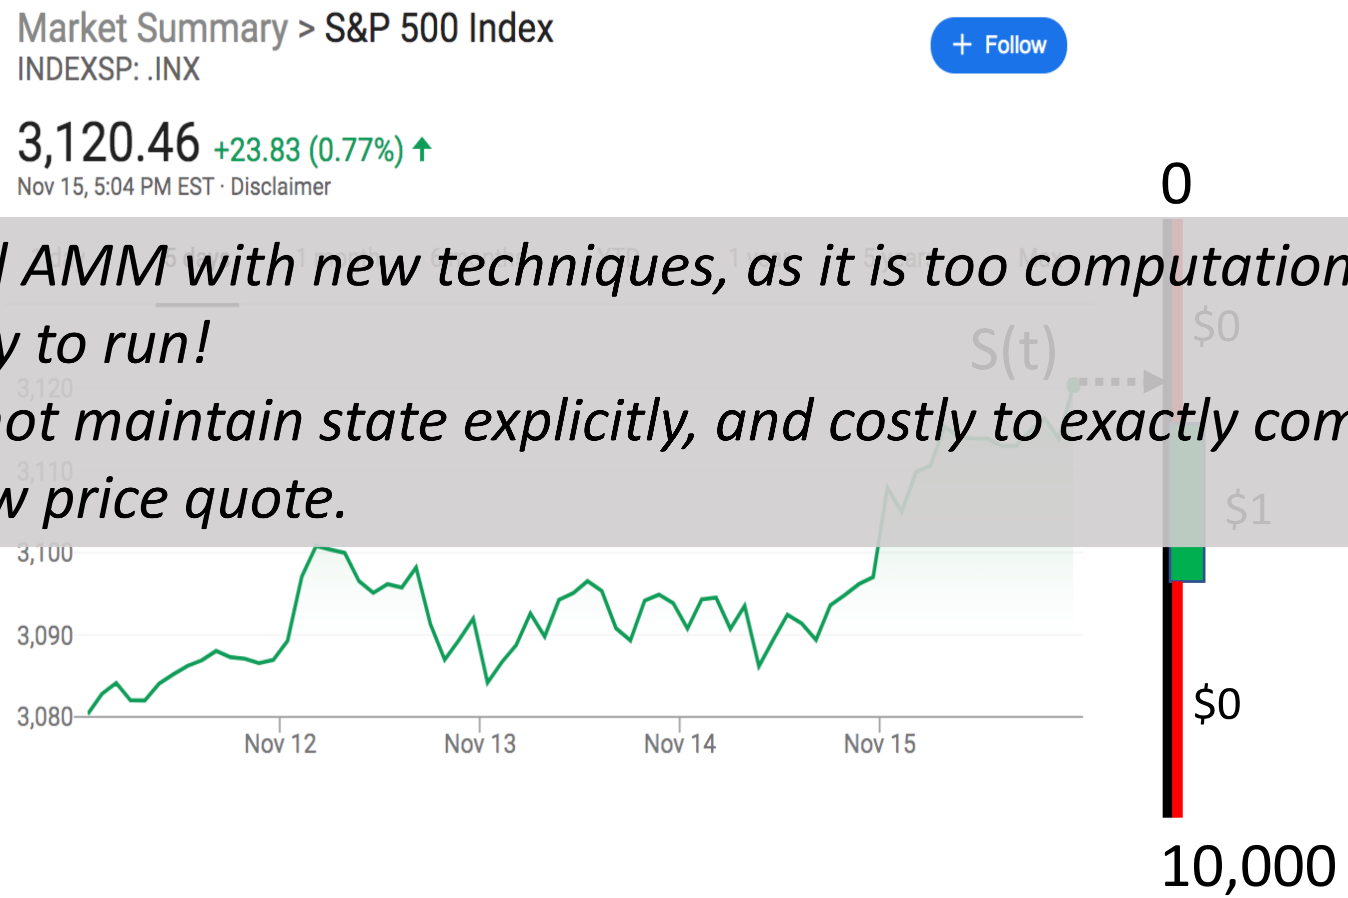

Need AMM with new techniques, as it is too computationally costly to run!
Cannot maintain state explicitly, and costly to exactly compute a new price quote.

Announcements

- HW2 will be out soon
- Office hours are extended (starting next week) to welcome more project discussions