A Brief Introduction to Prediction Markets

Jake Abernethy, University of Michigan

How do I find someone to bet with? Answer: A **Prediction Market**

- Prediction markets have existed for over 200 years.
- Typically, odds (prices) are set by supply and demand.
- People began to notice: the market prices are generally **very** accurate, and provide better predictors than expert assessments, etc. What's going on?
- Robin Hanson: "Rational expectations theory predicts that, in equilibrium, asset prices will reflect all of the information held by market participants. This theorized *information aggregation property* of prices has lead economists to become increasingly interested in using securities markets to predict future events."

Outline

1. Predictions Markets in Practice
2. Eliciting beliefs with proper scoring rules
3. Bregman divergences + proper scoring rules
4. Hanson's Market Scoring Rule
5. Securities markets

Example: Intrade

The Dark Knight Rises to break the all-time opening weekend box-office record

Last prediction was: **$4.08** / share

Today's Change: ▲ **+$0.02** (+0.5%)

Contract Type: **0-100** ?

40.8%
CHANCE

NASA to announce discovery of extraterrestrial life before midnight ET 31 Dec 2012

Last prediction was: **$0.43** / share

Today's Change: -

Contract Type: **0-100** ?

4.3%
CHANCE

The Most (in)Famous Prediction Market

- 1999: Intrade founded by John Delaney
- 2003: Acquired by TradeSports in 2003, later splits off after TradeSports closes down in 2008
- 2004: Intrade gains notoriety during Bush/Kerry election for providing continuous forecasts throughout campaign
- May 2011: Founder John Delaney dies at age 42 while climbing Mt. Everest, less than 50 meters from summit
- Nov. 2012: US regulator CFTC files suit against Intrade, leading Intrade to disallow US customers from betting
- Mar. 2013: Due to “financial irregularities”, Intrade halts trading, freezes all accounts. (Recently resolved.)

Iowa Electronic Markets (IEM): Legal and with Real Money

- Founded in 1988 at the University of Iowa for the purpose of research in market prediction accuracy
- Received a “no action” letter from the CFTC, permitting them to facilitate unregulated betting. (Such letters are apparently “no longer being given out”)
- On the downside, the IEM must obey a certain set of conditions. Most notably, individual traders may deposit no more than $500.

Predictious: A New *Bitcoin*-based Prediction Market

Example: Inkling Markets

Internal prediction markets used within companies

Intrade sells "Arrow-Debreu Securities"

Barack Obama to be re-elected President in 2012

Last prediction was: **$5.56** / share

Today's Change: ▲ **+$0.06** (+1.1%)

Contract Type: **0-100**

55.6%
CHANCE

"Arrow-Debreu Security": Contract pays \$10 if X happens, \$0 otherwise. If I think that Pr(X) = p then I should:

- Buy this security at any price less than \$10p
- Sell this security at any price greater than \$10p

Current price measures the population's collective beliefs

[1] Market Mechanisms: Continuous Double Auction (CDA)

- Used by Intrade.com and Betfair.com
- Market receives a sequence of **orders**
- **Two** types of orders:
  - **Limit** order: trader posts shares to order book
  - **Market** order: trader buys shares in order book

Predict | View All Un-Matched Predictions | Info | Rules

Best (highest) price members are buying at

Price per share	Quantity
$5.80	86 shares
$5.78	10 shares
$5.77	100 shares
$5.75	9 shares
$5.74	125 shares
$5.73	100 shares
$5.72	6 shares
$5.71	101 shares
$5.70	40 shares
$5.69	10 shares
$5.68	100 shares
$5.67	160 shares
$5.66	22 shares
$5.65	101 shares
$5.62	12 shares

Best (lowest) price members are selling at

Price per share	Quantity
$5.90	10 shares
$5.91	22 shares
$5.92	22 shares
$5.93	23 shares
$5.94	22 shares
$5.95	123 shares
$5.96	22 shares
$5.98	8 shares
$5.99	24 shares
$6.00	204 shares
$6.01	23 shares
$6.02	102 shares
$6.03	648 shares
$6.04	122 shares
$6.05	203 shares

Obama2012 Intrade: Bid+Ask+Trades

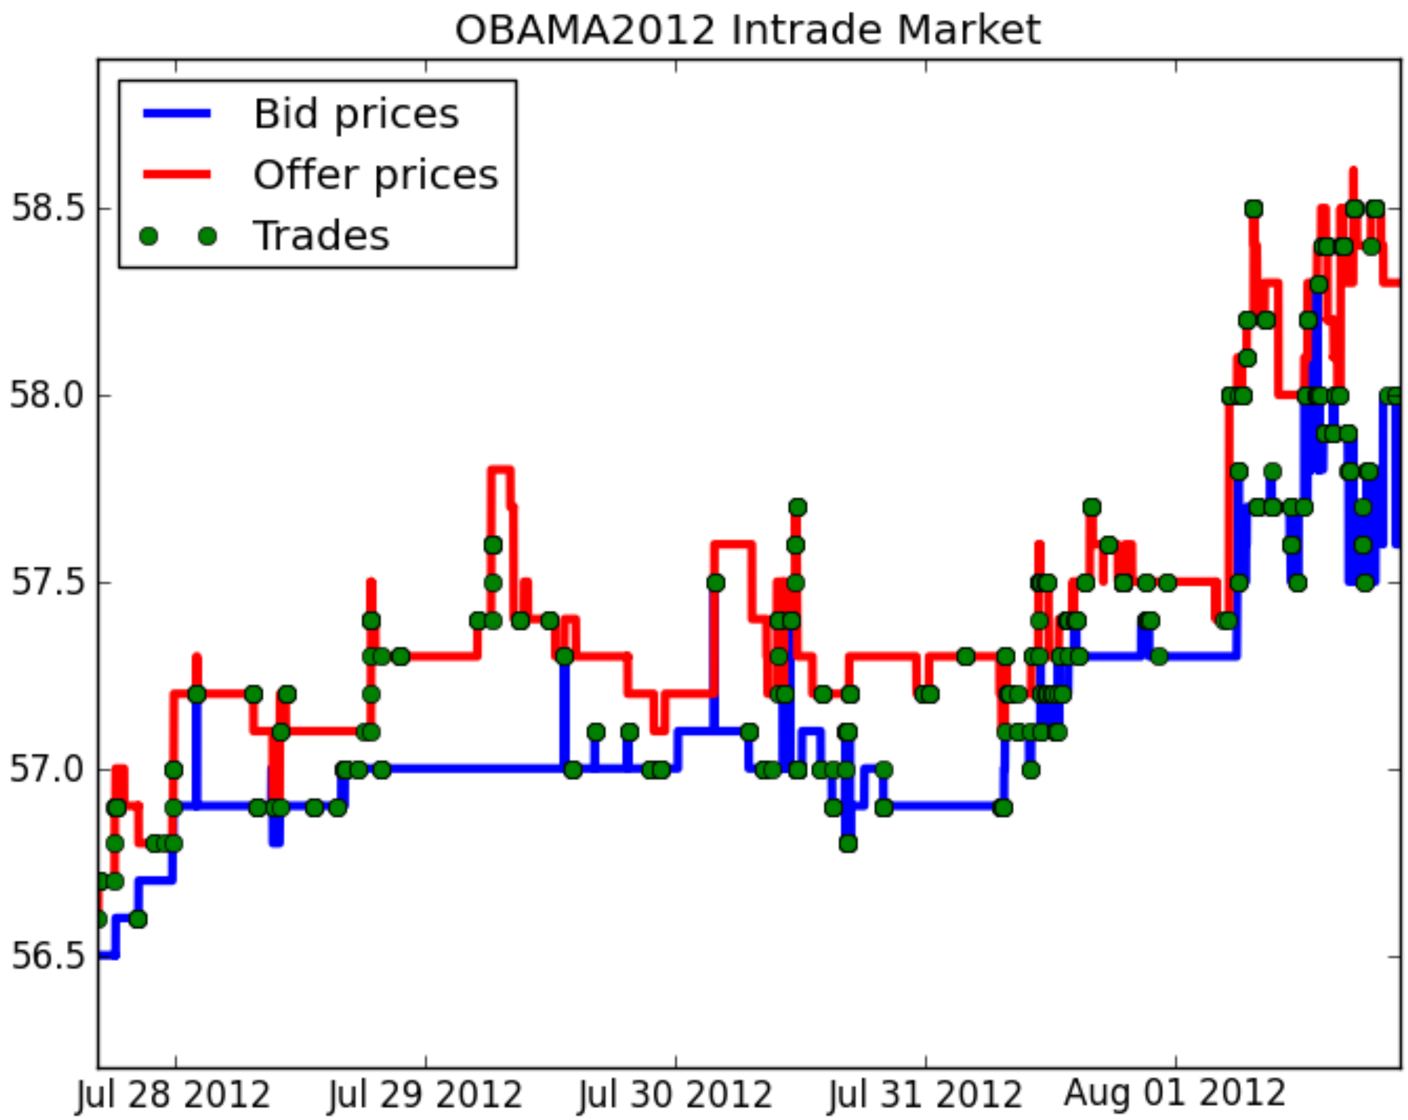

Aside: Problems with the CDA

- Chicken and egg problem: who is willing to join a market if there are no other participants?
- Not a lot of “liquidty”: it’s very easy to swing prices
- Large bid/ask spreads

- Alternative mechanism: the **automated market maker**, which we will be discussing later.

[2] How Quickly do Markets Respond?

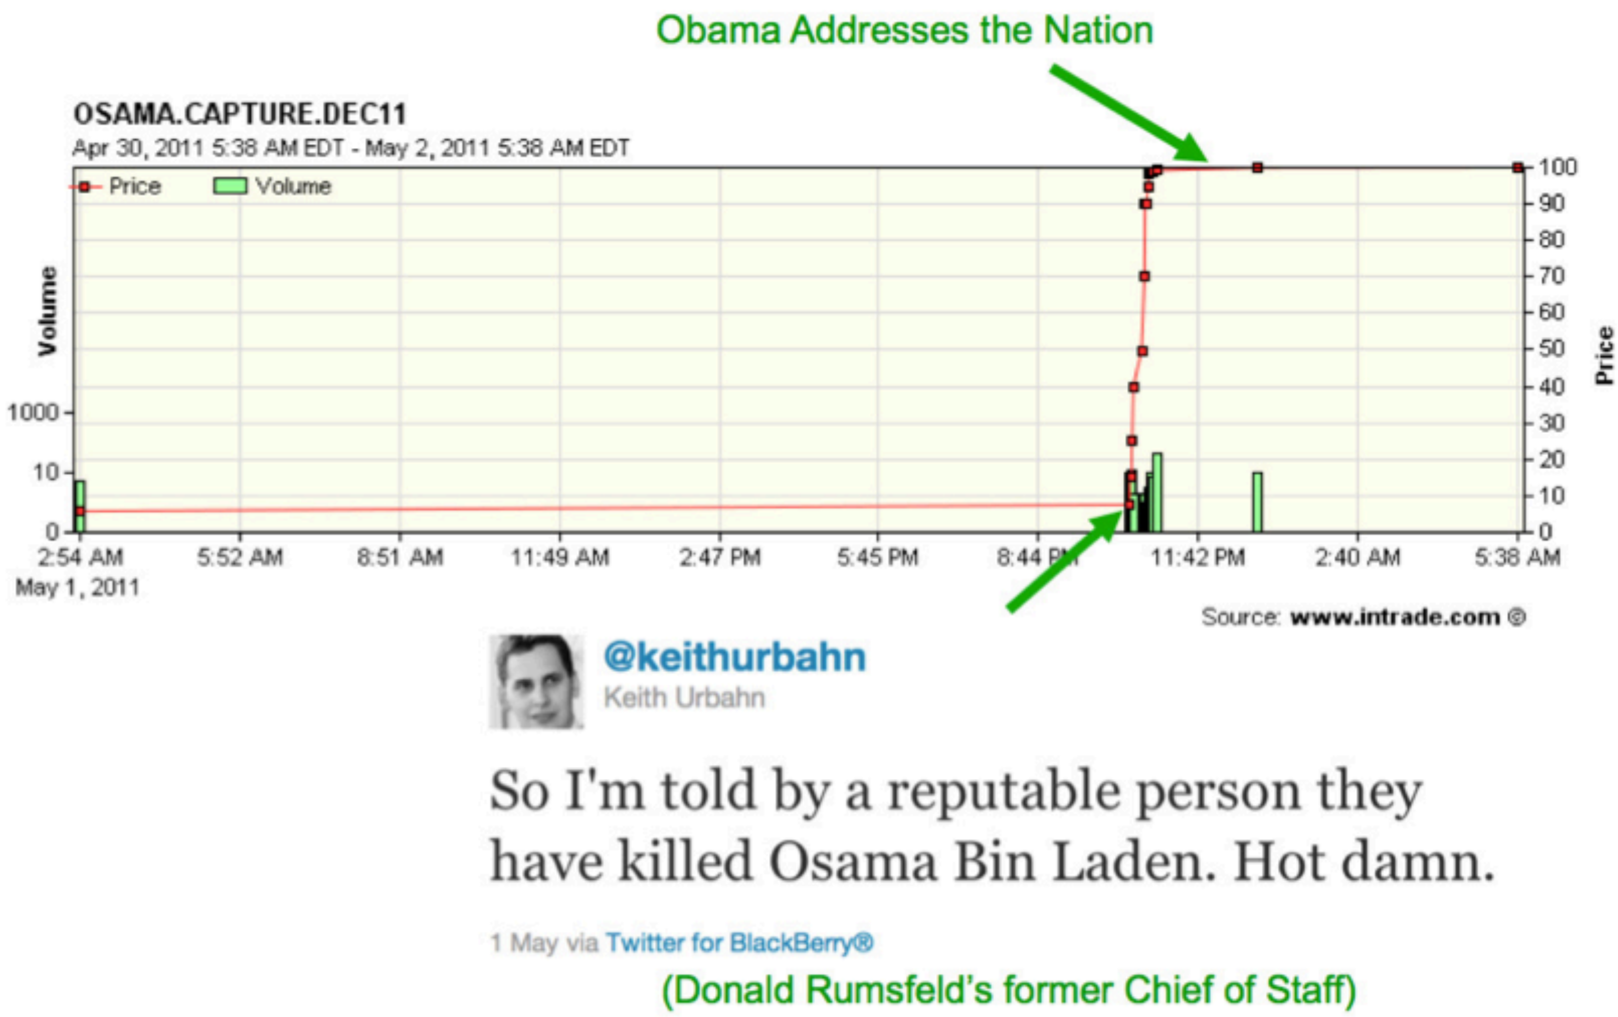

Source: Snowberg, Wolfers, Zitzewitz 2012

Market Prediction vs. True Vote Share

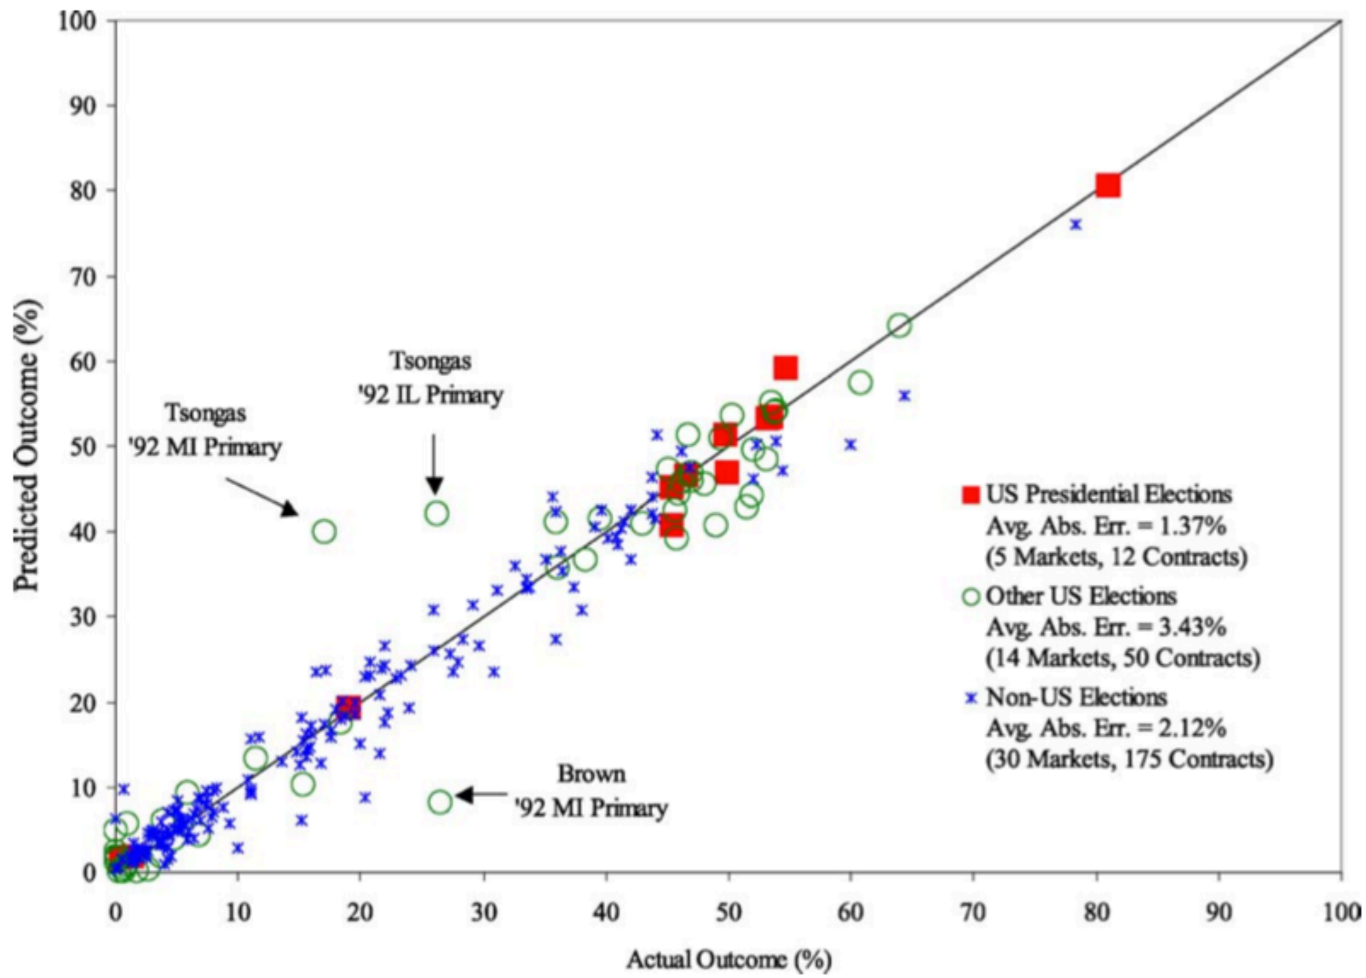

Berg et al., 2008: "Results From a Dozen Years of Election Futures Markets Research"

The Basics: Proper Scoring Rules

1950: Brier on Weather Forecasting

MONTHLY WEATHER REVIEW

EDITOR, JAMES E. CASKEY, JR.

Volume 78
Number 1

JANUARY 1950

Closed March 5, 1950
Issued April 15, 1950

VERIFICATION OF FORECASTS EXPRESSED IN TERMS OF PROBABILITY

GLENN W. BRIER

U. S. Weather Bureau, Washington, D. C.

[Manuscript received February 10, 1950]

INTRODUCTION

Verification of weather forecasts has been a controversial subject for more than a half century. There are a number of reasons why this problem has been so perplexing to meteorologists and others but one of the most important difficulties seems to be in reaching an agreement on the specification of a scale of goodness for weather forecasts. Numerous systems have been proposed but one of the greatest arguments raised against forecast verification is that forecasts which may be the "best" according to the accepted system of arbitrary scores may not be the most

numerically have been discussed previously [1, 2, 3, 4] so that the purpose here will not be to emphasize the enhanced usefulness of such forecasts but rather to point out how some aspects of the verification problem are simplified or solved.

VERIFICATION FORMULA

Suppose that on each of n occasions an event can occur in only one of r possible classes or categories and on one such occasion, i, the forecast probabilities are f_{i1}, f_{i2}, . . . f_{ir}, that the event will occur in classes 1, 2, . . . r, respectively. The r classes are chosen to be mutually

Table 2.—*Verification of a series of 85 forecasts expressed in terms of the probability of rain*

Forecast probability of rain	Observed proportion of rain cases	Forecast probability of rain	Observed proportion of rain cases
0.00–0.19	0.07	0.60–0.79	0.40
0.20–0.39	.10	0.80–1.00	.50
0.40–0.59	.29		

How Should We Pay a Forecaster?

What is the "right" payment scheme to reward/punish a forecaster who makes a sequence of probability predictions for events that we observe?

- The sequence of outcomes: $y_1, y_2, y_3, \ldots \in \{0,1\}$
- The sequence of forecasts: $p_1, p_2, p_3, \ldots \in [0,1]$
- The forecaster's payment: $\frac{1}{T}\sum_{t=1}^{T} S(y_t, p_t)$

Brier Score ⇔ Quadratic Scoring Rule

- For a binary outcome $y \in \{0,1\}$, $p \in [0,1]$

$$S(y,p) = -(y-p)^2$$

- For one of n outcomes, $y \in \{1,\ldots,n\}$, $\mathbf{p} \in \Delta_n$

$$S(y,\mathbf{p}) = -\sum_{i=1}^{n} (\mathbf{1}_{y=i} - p_i)^2$$

What' s Special About This Function?

$$S(y,p) = -(y-p)^2$$

Assume y is random and $\Pr(y=1)=q$. Then...

$$\operatorname*{argmax}_{p\in[0,1]}\left(\mathrm{E}\left[-(y-p)^2\right]\right)$$

$$= \operatorname*{argmax}_{p\in[0,1]}\left(-q(1-p)^2-(1-q)p^2\right)$$

$$= \operatorname*{argmax}_{p\in[0,1]}\left(-(p-q)^2-q+q^2\right)=q$$

Proper Scoring Rules

- What we have just introduced is the notion of a proper scoring rule, any function S satisfying

$$\mathrm{E}_{y\sim\mathbf{q}}\left[S(y,\mathbf{q})\right] \geq \mathrm{E}_{y\sim\mathbf{q}}\left[S(y,\mathbf{p})\right] \quad \forall \mathbf{p},\mathbf{q} \in \Delta_n$$

- The scoring rule is said to be strictly proper if the above inequality is strict unless $\mathbf{p} = \mathbf{q}$

Another Strictly Proper Scoring Rule

$$S(y,\mathbf{p}) = \log p(y)$$

- This is known as the logarithmic scoring rule. For binary random variables, it can be written as:

$$S(y,p) = \begin{cases} \log p & y=1 \\ \log(1-p) & y=0 \end{cases}$$

- **EXERCISE:** check that this is proper!

Digression: Bregman Divergences

- A Bregman divergence measures distance with respect to a convex function f

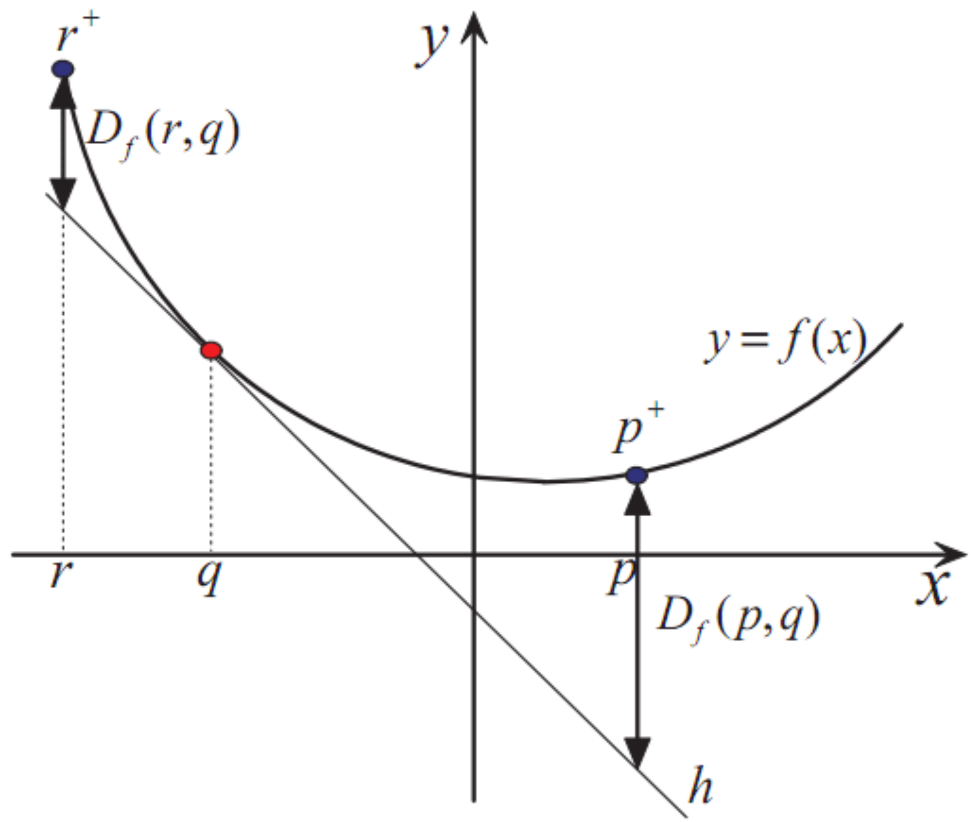

$$D_f(p,q) = f(p) - f(q) - \nabla f(q) \cdot (p - q)$$

Digression: Bregman Divergences

$$D_f(p,q) = f(p) - f(q) - \nabla f(q) \cdot (p - q)$$

- **Properties:**

$$D_f(p,p) = 0 \quad \forall p$$

$$D_f(p,q) \geq 0 \quad \forall p,q$$

$$D_f(p,q) \neq D_f(q,p) \quad (in \ \ general)$$

Bregman Divergences III

$$D_f(p,q) = f(p) - f(q) - \nabla f(q) \cdot (p - q)$$

- **Example 1**, quadratic:

$$f(p) = \| p \|^2 \Rightarrow D_f(p,q) = \| p - q \|^2$$

- **Example 2**, entropic:

$$f(\mathbf{p}) = \sum_i p_i \log p_i \Rightarrow D_f(\mathbf{p},\mathbf{q}) = \sum_i p_i \log \frac{p_i}{q_i}$$

Bregman Diverg.⇔ Scoring Rule??

- Let f be any convex function
- Let $\mathbf{e}_i$ be the ith indicator vector, $\mathbf{e}_i = \langle 0,\ldots,0,1,0,\ldots,0\rangle$
- Let **p, q** be any two distributions
- Fact: There exists a function g such that

$$\mathrm{E}_{i\sim\mathbf{q}}\left[D_f(\mathbf{e}_i,\mathbf{p})\right] = D_f(\mathbf{q},\mathbf{p}) + g(\mathbf{q})$$

and so

$$\arg\max_{\mathbf{p}\in\Delta_n} \mathrm{E}_{i\sim\mathbf{q}}\left[-D_f(\mathbf{e}_i,\mathbf{p})\right] = \mathbf{q}$$

- This is the scoring rule property!!

Bregman Diverg. ⇔ Scoring Rule!!

- We now have a recipe for constructing scoring rules: Take any convex function f and set

$$S(i,\mathbf{p}) = -D_f(\mathbf{e}_i,\mathbf{p})$$

- Quadratic Scoring Rule: $f(\mathbf{p}) = \|\mathbf{p}\|_2^2$

- Log Scoring Rule: $f(\mathbf{p}) = \sum_i p_i \log p_i$

Market Scoring Rules

Robin Hanson proposed the following idea to create a prediction market based on an automated market maker:

- Suppose we have a random variable X which will take one of n values $\{1, 2, \ldots, n\}$
- The MM chooses a scoring rule S and announces it
- The MM then posts an initial distribution (prior) $\mathbf{p}_0$
- Traders arrive, one-by-one, giving updates $\mathbf{p}_{t-1} \to \mathbf{p}_t$
- Eventually, outcome X is revealed, and trader t earns (or loses)

$$S(X, \mathbf{p}_t) - S(X, \mathbf{p}_{t-1})$$

Incentives and Costs

- Assume trader t has belief distribution $\mathbf{p}$ on X, which can (and should!) depend on previous market observations
- Suppose he wants to maximize his payment

$$\operatorname{argmax}_{\mathbf{p}_t} \mathrm{E}_{X\sim\mathbf{p}}[S(X,\mathbf{p}_t) - S(X,\mathbf{p}_{t-1})]$$

$$= \operatorname{argmax}_{\mathbf{p}_t} \mathrm{E}_{X\sim\mathbf{p}}[S(X,\mathbf{p}_t)] = \mathbf{p}$$

- The MM must make all payments, which total

$$\sum_{t=1}^{T}\left[S(X,\mathbf{p}_t) - S(X,\mathbf{p}_{t-1})\right] = S(X,\mathbf{p}_T) - S(X,\mathbf{p}_0)$$

- This is bounded! This is like MM's subsidy to market.

LMSR: Log Market Scoring Rule

- Initial hypothesis $\mathbf{p}_0$ is the uniform distribution
- Trader t posts an update $\mathbf{p}_{t\text{-}1} \rightarrow \mathbf{p}_t$
- After X is revealed, trader t earns $\log(\mathbf{p}_t(X)/\mathbf{p}_{t\text{-}1}(X))$

- Hanson: the LMSR is an important special case, the only MSR for which "betting on conditional probabilities does not affect marginal probabilities"

- The market maker's worst case loss is bounded by $\log n$, where n is the number of possible values of X

Arrow–Debreu Securities

Barack Obama to be re-elected President in 2012

Last prediction was: **$5.56** / share

Today's Change: ▲ **+$0.06** (+1.1%)

Contract Type: **0-100**

55.6%
CHANCE

Arrow–Debreu Securities

Potential payoff is \$10. If I think that the probability of this event is p, I should

- Buy this security at any price less than \$10p
- Sell this security at any price greater than \$10p

Current price measures the population's collective beliefs

Arrow–Debreu Securities

Barack Obama to be re-elected President in 2012

Last prediction was **$5.56** / share → **55.6%** CHANCE

Today's Change: ▲ **+$0.06** (+1.1%)

Contract Type: **0-100**

How do we arrive at the current price?

Arrow–Debreu Securities

How do we arrive at the current price?

- Traditional stock market style pricing (continuous double auction) – low liquidity, huge spreads

Arrow–Debreu Securities

How do we arrive at the current price?

- Traditional stock market style pricing (continuous double auction) – low liquidity, huge spreads
- Automated market maker – willing to risk a (bounded) loss in order to encourage trades

Market Makers for Complete Markets

- In a complete market, a security is offered for each of a set of mutually exclusive and exhaustive events

Market Makers for Complete Markets

- In a complete market, a security is offered for each of a set of mutually exclusive and exhaustive events

Worth
$1 iff

Worth
$1 iff

Market Makers for Complete Markets

- In a complete market, a security is offered for each of a set of mutually exclusive and exhaustive events

Worth $1 iff

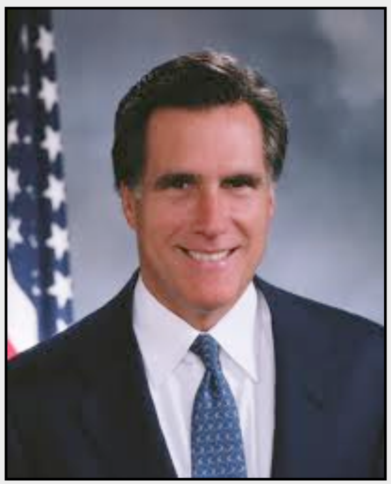

Worth $1 iff

- An automated market maker is always willing to buy and sell these securities at some price

Cost Functions

Cost Functions

Worth $1 iff

Worth $1 iff

Already purchased: q_1 shares q_2 shares

Cost Functions

Worth
$1 iff

Worth
$1 iff

Already purchased:	q_1 shares	q_2 shares
Want to purchase:	r_1 shares	r_2 shares

Cost Functions

Cost of purchase:

$$C(\mathbf{q} + \mathbf{r}) - C(\mathbf{q})$$

Worth $1 iff

Worth $1 iff

Already purchased:	q_1 shares	q_2 shares
Want to purchase:	r_1 shares	r_2 shares

Cost Functions

Cost of purchase:

$$C(\mathbf{q} + \mathbf{r}) - C(\mathbf{q})$$

Worth $1 iff

Worth $1 iff

Already purchased:	q_1 shares	q_2 shares
Want to purchase:	r_1 shares	r_2 shares
Instantaneous prices:	$p_1 = \partial C / \partial q_1$	$p_2 = \partial C / \partial q_2$

Cost Functions

Cost of purchase:

$$C(\mathbf{q}+\mathbf{r}) - C(\mathbf{q})$$

Worth $1 iff

Worth $1 iff

Already purchased:	q_1 shares	q_2 shares
Want to purchase:	r_1 shares	r_2 shares
Instantaneous prices:	$p_1 = \partial C / \partial q_1$	$p_2 = \partial C / \partial q_2$

"predictions"

Back to the LMSR

Remember the logarithmic market scoring rule…

- Initial hypothesis $\mathbf{p}_0$ is the uniform distribution

- Trader t posts an update $\mathbf{p}_{t\text{-}1} \rightarrow \mathbf{p}_t$

- After outcome i is revealed, trader t receives

$$\log(p_{t,i}) - \log(p_{t\text{-}1,i}) = \log(p_{t,i} / p_{t\text{-}1,i})$$

Back to the LMSR

The logarithmic market scoring rule can be implemented as a cost function based market with cost function

$$C(q_1,\ldots,q_N) = \log \sum_{i=1}^{N} \exp(q_i)$$

and instantaneous prices

$$p_i = \frac{\exp(q_i)}{\Sigma_j \exp(q_j)}$$

Back to the LMSR

The logarithmic market scoring rule can be implemented as a cost function based market with cost function

$$C(q_1,\ldots,q_N) = \log \sum_{i=1}^{N} \exp(q_i)$$

and instantaneous prices

$$p_i = \frac{\exp(q_i)}{\Sigma_j \exp(q_j)}$$

Notice that p_i is increasing in q_i and the prices sum to 1

THANK YOU!