Interpreting Prediction Markets: a Stochastic Approach

Rafael Frongillo

Department of Computer Science
University of California at Berkeley

July 1, 2012

Joint work with Nicolás Della Penna and Mark Reid

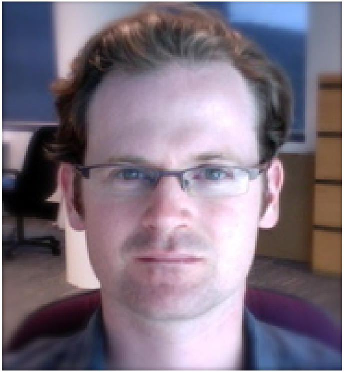

Work done while visiting ANU + NICTA

Standard Prediction Market

Obama $1		Romney $1		
$0.54	+	$0.46	=	1

- Traders buy and sell contracts
- Prices fluctuate as demand changes
- Prices should reflect "consensus estimate"

What is this

Standard Prediction Market

Obama $1	Romney $1
$0.54	$0.46

$0.54 + $0.46 = 1

- Traders buy and sell contracts
- Prices fluctuate as demand changes
- Prices should reflect “consensus estimate”

What is this

Standard Prediction Market

Obama		Romney		
$1		$1		
$0.54	+	$0.46	=	1

- Traders buy and sell contracts
- Prices fluctuate as demand changes
- Prices should reflect “consensus estimate”

What is this

Standard Prediction Market

Obama	Romney		
\$1	\$1		
\$0.54 +	\$0.46	=	1

- Traders buy and sell contracts
- Prices fluctuate as demand changes
- Prices should reflect “consensus estimate”

What is this

Standard Prediction Market

Obama	Romney
$1	$1

\$0.54 + \$0.46 = 1

- Traders buy and sell contracts
- Prices fluctuate as demand changes
- Prices should reflect “consensus estimate”

What is this

Standard Prediction Market

Obama $1		Romney $1		
$0.54	+	$0.46	=	1

- Traders buy and sell contracts
- Prices fluctuate as demand changes
- Prices should reflect "consensus estimate"

What is this?

Standard Prediction Market

Obama $1		Romney $1		
$0.54	+	$0.46	=	1

- Traders buy and sell contracts
- Prices fluctuate as demand changes
- Prices should reflect "consensus estimate"

What is this
in terms of the *traders' beliefs*?

Answer from Standard Theory

If traders have unbounded wealth and are risk neutral,
prices = **last traders' belief**

If traders perform proper Bayesian updating,
prices = **posterior** given everyone's private info

Big If's!!

Answer from Standard Theory

If traders have unbounded wealth and are risk neutral,
prices = **last traders' belief**

If traders perform proper Bayesian updating,
prices = **posterior** given everyone's private info

Big If's!!

Answer from Standard Theory

If traders have unbounded wealth and are risk neutral,
prices = **last traders' belief**

If traders perform proper Bayesian updating,
prices = **posterior** given everyone's private info

Big If's!!

Standard Equilibrium Analysis

Setting:

- Look at the *distribution* $\mathcal{P}$ of traders' beliefs
- Fix some price π *of contract 1, say*
- Look at total demand for that price
- Equilibrium is π^* s.t. supply = demand:

$$\int_0^1 \text{demand}(\pi^*, p)\, d\mathcal{P}(p) = 0$$

Note: demand for Obama = supply for Romney

Standard Equilibrium Analysis

Setting:

- Look at the *distribution* $\mathcal{P}$ of traders' beliefs
- Fix some price π *of contract 1, say*
- Look at total demand for that price
- Equilibrium is π^* s.t. supply = demand:

$$\int_0^1 \text{demand}(\pi^*, p)\, d\mathcal{P}(p) = 0$$

Note: demand for Obama = supply for Romney

Standard Equilibrium Analysis: Results

Manski, 2004:

- Risk neutral traders $\implies$ π^* = quantile of $\mathcal{P}$
- Equlibrium point ("Manski point"): π^* such that

$$\frac{\int_0^{\pi^*} \mathcal{P}(p)dp}{1-\pi^*} = \frac{\int_{\pi^*}^{1} \mathcal{P}(p)dp}{\pi^*} \implies \int_{\pi^*}^{1} \mathcal{P}(p)dp = \pi^*$$

Wolfers and Zitzewitz, 2006:

- Kelly bettor: demand = $\frac{W}{\pi}\frac{p-\pi}{1-\pi}$ *linear in p*
- Kelly bettors $\implies$ π^* = mean of $\mathcal{P}$

Standard Equilibrium Analysis: Results

Manski, 2004:

- Risk neutral traders $\Longrightarrow$ π^* = quantile of $\mathcal{P}$
- Equlibrium point (“Manski point”): π^* such that

$$\frac{\int_0^{\pi^*} \mathcal{P}(p)dp}{1-\pi^*} = \frac{\int_{\pi^*}^1 \mathcal{P}(p)dp}{\pi^*} \Longrightarrow \int_{\pi^*}^1 \mathcal{P}(p)dp = \pi^*$$

Wolfers and Zitzewitz, 2006:

- Kelly bettor: demand = $\dfrac{W}{\pi} \dfrac{p-\pi}{1-\pi}$ *linear in p*
- Kelly bettors $\Longrightarrow$ π^* = mean of $\mathcal{P}$

Standard Equilibrium Analysis: Really?

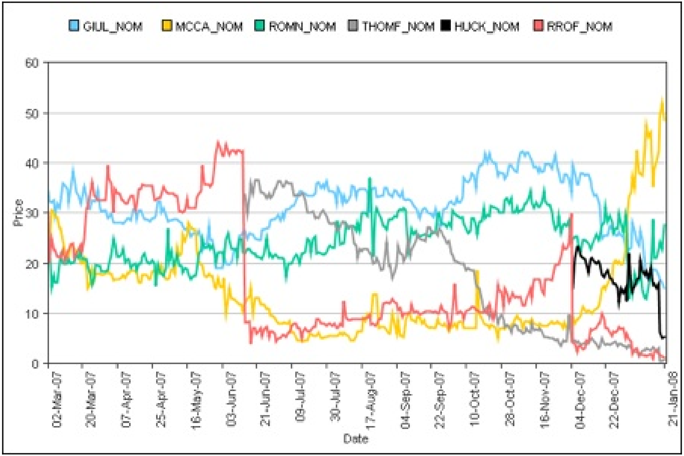

Where is the equilibrium?

How do we use these prices to make predictions?

Standard Equilibrium Analysis: Really?

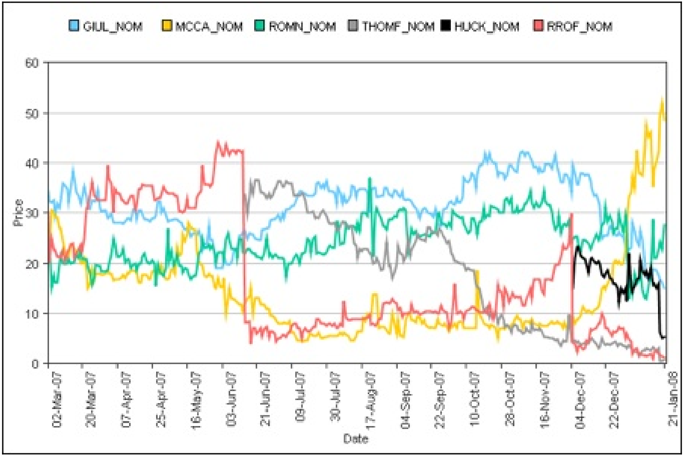

Where is the equilibrium?

How do we use these prices to make predictions?

Standard Equilibrium Analysis: Really?

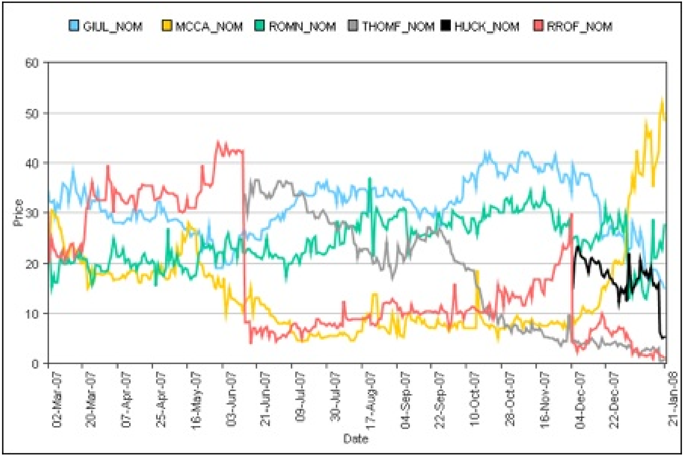

Where is the equilibrium?

How do we use these prices to make predictions?

A Stochastic Approach

Based on Othman and Sandholm, 2010:

- Look at a *sequential* market model
- *Sample* traders repeatedly from $\mathcal{P}$
- Each trades one-by-one in the market
- The current prices *adjust* to trades

  ... using an automated market maker

The Market Maker

Use model of Abernethy, Chen, Vaughan (2011)

- n mutually exclusive events
- Convex function $C : \mathbb{R}^n \to \mathbb{R}$
- Current quantity $q \in \mathbb{R}^n$
- Trade of $d \in \mathbb{R}^n$ costs $C(q+d) - C(q)$
- Prices: $\nabla C(q)$

Example: exponential weights

$$C(q) = \log\left(\sum_i \exp(q_i)\right) \qquad \nabla C(q_i) = \frac{\exp(q_i)}{\sum_j \exp(q_j)}$$

The Market Maker

Use model of Abernethy, Chen, Vaughan (2011)

- n mutually exclusive events
- Convex function $C : \mathbb{R}^n \to \mathbb{R}$
- Current quantity $q \in \mathbb{R}^n$
- Trade of $d \in \mathbb{R}^n$ costs $C(q + d) - C(q)$
- Prices: $\nabla C(q)$

Example: exponential weights

$$C(q) = \log\left(\sum_i \exp(q_i)\right) \qquad \nabla C(q_i) = \frac{\exp(q_i)}{\sum_j \exp(q_j)}$$

The Demands

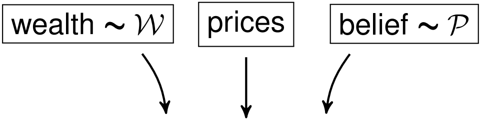

$$d\ (W,\quad \pi,\quad p)\ \in \mathbb{R}^n$$

d_i is the demand for contract i

The Full Model

For $t = 1 \ldots T$:

- Draw a trader (p, W) from $(\mathcal{P}, \mathcal{W})$ *i.i.d.*
- Trader buys bundle $d(W, \pi_t, p)$
- Price adjusts:

$$\pi_{t+1} \leftarrow \nabla C\Big((\nabla C)^{-1}(\pi_t) + d(W, \pi_t, p)\Big)$$

The Stationary Point

First question: what is the "fixed point" of our process?

When does $\mathbb{E}[\pi_{t+1}] = \pi_t$?

Define π^s to be this stationary point

The Stationary Point

First question: what is the “fixed point” of our process?

When does $\mathbb{E}[\pi_{t+1}] = \pi_t$?

Define π^s to be this stationary point

The Stationary Point

First question: what is the “fixed point” of our process?

When does $\mathbb{E}[\pi_{t+1}] = \pi_t$?

Define π^s to be this stationary point

Stationarity and Equilibrium

How does this π^s relate to the equilibrium π^*?

Othman and Sandholm:

$$\left.\begin{array}{l} \text{■ Binary market} \\ \text{■ Risk neutral traders} \\ \text{■ Each invest } \epsilon \text{ at a time} \end{array}\right\} \implies \pi^s = \pi^*$$

Stationarity and Equilibrium: Our Model

Is this a general phenomenon?

Does $\pi^s = \pi^*$ in our more general model?

Theorem 1

For very general demands d, $\pi^s \to \pi^*$ as $W \to 0$

Obtain Othman and Sandholm result as a corollary

Stationarity and Equilibrium: Our Model

Is this a general phenomenon?

Does $\pi^s = \pi^*$ in our more general model?

Theorem 1

For very general demands d, $\pi^s \to \pi^*$ as $W \to 0$

Obtain Othman and Sandholm result as a corollary

Stationarity and Equilibrium: Our Model

Is this a general phenomenon?

Does $\pi^s = \pi^*$ in our more general model?

Theorem 1

For very general demands d, $\pi^s \to \pi^*$ as $W \to 0$

Obtain Othman and Sandholm result as a corollary

Fixed vs. Continuous Prices

Same demands for both equilibrium and market-maker settings

fixed continuous

Open question: does Theorem 1 hold with more sensible demands?

Fixed vs. Continuous Prices

Same demands for both equilibrium and market-maker settings

fixed → equilibrium

continuous → market-maker

Open question: does Theorem 1 hold with more sensible demands?

Fixed vs. Continuous Prices

Same demands for both equilibrium and market-maker settings

equilibrium → fixed

market-maker → continuous

Open question: does Theorem 1 hold with more sensible demands?

Market Making and Online Learning

Chen and Vaughan, 2009: Market maker update = FTRL

Follow the Regularized Leader:

- Losses $\ell_t \in \mathbb{R}^n$
- Actions $\mathbf{w}_t \in \Delta_n$
- Convex regularizer R
- $\mathbf{w}_{t+1} \leftarrow \operatorname{argmin}\left\{\sum \ell_t \cdot \mathbf{w} + R(\mathbf{w})\right\}$

Regret:

$$\sum_t \ell_t \cdot \mathbf{w}_t - \min_i \left(\sum_t \ell_t\right)_i$$

Market Making and Online Learning

Chen and Vaughan, 2009: Market maker update = FTRL

Follow the Regularized Leader:

- Losses $\ell_t \in \mathbb{R}^n$
- Actions $\mathbf{w}_t \in \Delta_n$
- Convex regularizer R
- $\mathbf{w}_{t+1} \leftarrow \operatorname{argmin}\left\{\sum \ell_t \cdot \mathbf{w} + R(\mathbf{w})\right\}$

Regret:

$$\sum_t \ell_t \cdot \mathbf{w}_t - \min_i \left(\sum_t \ell_t\right)_i$$

Matching Up the Final Losses

Important caveat: final loss terms do not match up completely

FTRL regret	$\sum_t \ell_t \cdot \mathbf{w}_t$	$-$	$\min_i \left(\sum_t \ell_t\right)_i$
Market Maker gain	$C\left(\sum_t d_t\right) - C(0)$	$-$	$\max_i \left(\sum_t d_t\right)_i$

$$\sum_t d_t \cdot \pi_t$$

Observation: They *do* line up in the fixed price model! $d_t = -\ell_t$

Matching Up the Final Losses

Important caveat: final loss terms do not match up completely

FTRL regret $\quad \sum_t \ell_t \cdot \mathbf{w}_t \quad - \quad \min_i \left(\sum_t \ell_t \right)_i$

Market Maker gain $\quad C\left(\sum_t d_t \right) - C(0) \quad - \quad \max_i \left(\sum_t d_t \right)_i$

$$\sum_t d_t \cdot \pi_t$$

Observation: They *do* line up in the fixed price model! $\quad d_t = -\ell_t$

FTRL → Mirror Descent

Observation: If losses are *gradients*, FTRL = Mirror Descent

$$\ell_{t+1} = \nabla f(\mathbf{w}_t)$$

And *stochastic* gradients → *Stochastic* Mirror Descent (SMD)

$$\ell_{t+1} = \nabla F(\mathbf{w}_t; \xi)$$

Theorem 2

If demands are (negative) gradients

$$d(W, \pi, p) = -\nabla F(\pi; p, W),$$

our stochastic update is an SMD of

$$f(\pi) := \mathbb{E}[F(\pi; W, p)].$$

FTRL → Mirror Descent

Observation: If losses are *gradients*, FTRL = Mirror Descent

$$\ell_{t+1} = \nabla f(\mathbf{w}_t)$$

And *stochastic* gradients → *Stochastic* Mirror Descent (SMD)

$$\ell_{t+1} = \nabla F(\mathbf{w}_t; \xi)$$

Theorem 2

If demands are (negative) gradients

$$d(W, \pi, p) = -\nabla F(\pi; p, W),$$

our stochastic update is an SMD of

$$f(\pi) := \mathbb{E}[F(\pi; W, p)].$$

FTRL $\rightarrow$ Mirror Descent

Observation: If losses are *gradients*, FTRL = Mirror Descent

$$\ell_{t+1} = \nabla f(\mathbf{w}_t)$$

And *stochastic* gradients $\rightarrow$ *Stochastic* Mirror Descent (SMD)

$$\ell_{t+1} = \nabla F(\mathbf{w}_t; \xi)$$

Theorem 2

If demands are (negative) gradients

$$d(W, \pi, p) = -\nabla F(\pi; p, W),$$

our stochastic update is an SMD of

$$f(\pi) := \mathbb{E}[F(\pi; W, p)].$$

Application: Kelly Bettors

Back to Kelly bettors: $$d(W,\pi,p) = \frac{W}{\pi}\frac{p-\pi}{1-\pi}$$

(note: fixed-price model)

Can write d as a gradient of $F(\pi; W, p) := W \cdot \mathrm{KL}(p, \pi)$

Corollary

The stochastic model with Kelly bettors is an SMD of

$$f(\pi) := \overline{W} \cdot KL(\overline{p}, \pi)$$

Note: generalizes Wolfers and Zitzewitz by Theorem 1!

Application: Kelly Bettors

Back to Kelly bettors: $$d(W, \pi, p) = \frac{W}{\pi} \frac{p - \pi}{1 - \pi}$$

(note: fixed-price model)

Can write d as a gradient of $$F(\pi; W, p) := W \cdot \mathrm{KL}(p, \pi)$$

Corollary

The stochastic model with Kelly bettors is an SMD of

$$f(\pi) := \overline{W} \cdot KL(\overline{p}, \pi)$$

Note: generalizes Wolfers and Zitzewitz by Theorem 1!

Application: Kelly Bettors

Back to Kelly bettors: $$d(W,\pi,p) = \frac{W}{\pi}\,\frac{p-\pi}{1-\pi}$$

(note: fixed-price model)

Can write d as a gradient of $$F(\pi; W, p) := W \cdot \mathrm{KL}(p,\pi)$$

Corollary

The stochastic model with Kelly bettors is an SMD of

$$f(\pi) := \overline{W} \cdot KL(\overline{p},\pi)$$

Note: generalizes Wolfers and Zitzewitz by Theorem 1!

Back to Interpreting Prediction Markets

SMD has optimization guarantees:

Theorem (Duchi, Shalev-Shwartz, Singer, Tewari)

If $\|\nabla F(\pi; p)\|^2 \le G^2$ *for all* p, π, *and* R *is* σ*-strongly convex, then with probability* $1 - \delta$,

$$f(\overline{\pi}_T) \le \min_{\pi} f(\pi) + \left(\frac{D^2}{\eta T} + \frac{G^2 \eta}{2\sigma}\right)\left(1 + 4\sqrt{\log\frac{1}{\delta}}\right).$$

Time-averaged price!

Back to Interpreting Prediction Markets

SMD has optimization guarantees:

Theorem (Duchi, Shalev-Shwartz, Singer, Tewari)

If $\|\nabla F(\pi; p)\|^2 \leq G^2$ *for all* p, π, *and* R *is* σ*-strongly convex, then with probability* $1-\delta$,

$$f(\overline{\pi}_T) \leq \min_{\pi} f(\pi) + \left(\frac{D^2}{\eta T} + \frac{G^2 \eta}{2\sigma}\right)\left(1 + 4\sqrt{\log\frac{1}{\delta}}\right).$$

Time-averaged price!

Average the Prices

So perhaps one should *average* market prices to form predictions

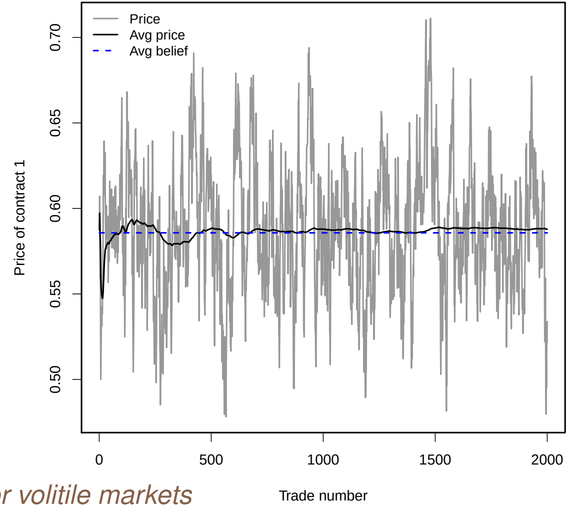

Particularly for volitile markets

Future Directions

- Sensible demands for stationarity (Theorem 1)
- Using results on *learning rates* to set market *liquidity*
- Try on real market data!

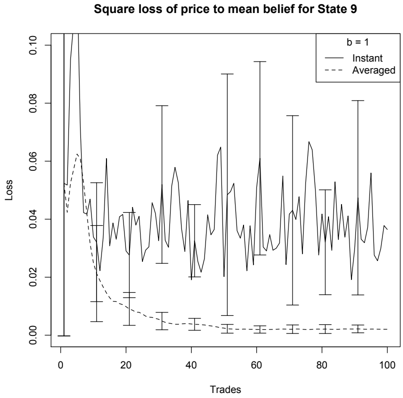

thank you