General Characterizations of Truthfulness via Convex Analysis

Rafael Frongillo

Department of Computer Science
University of California at Berkeley

November 29, 2012

Joint work with Ian Kash (MSRC)

Ian Says Hi!

Warm-up: Convex Functions

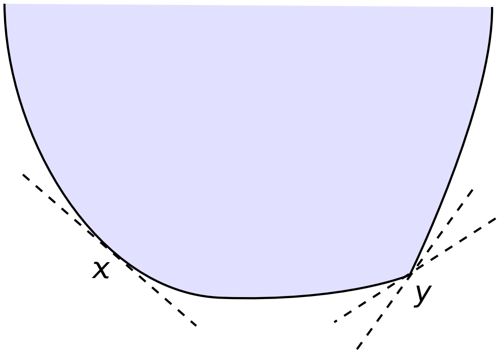

Definition

$G : \mathcal{T} \to \mathbb{R}$ is *convex* if for all $x, y \in \mathcal{T}$ and all $\alpha \in [0, 1]$

$$\alpha G(x) + (1 - \alpha)G(y) \geq G(\alpha x + (1 - \alpha)y)$$

Warm-up: Convex Functions

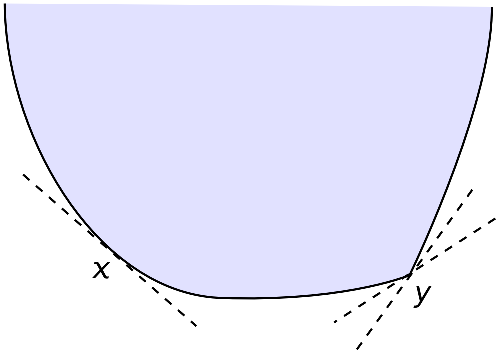

Definition

A linear function $dG_t : \mathcal{T} \to \mathbb{R}$ is a *subgradient* to G at t if

$$\forall t' \in \mathcal{T} \quad G(t') \geq G(t) + dG_t(t' - t)$$

Pointwise Supremum

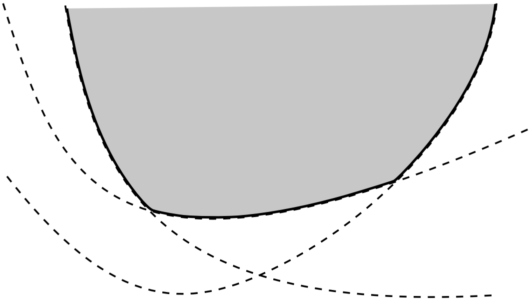

Fact

If G_i are convex functions for $i \in I$, then G is convex:

$$G(t) := \sup_{i \in I} G_i(t)$$

Mechanism Design

Single-player mechanism:

- Outcome space $\mathcal{O}$ *possible allocations*
- Type space $\mathcal{T} = (\mathcal{O} \to \mathbb{R})$ *valuation functions*
- Allocation rule $a : \mathcal{T} \to \mathcal{O}$ *reports to outcomes*
- Payment rule $p : \mathcal{T} \to \mathbb{R}$ *reports to payments*

Bidder with type t who reports $t' \in \mathcal{T}$ has net utility

$$U(t', t) = t(a(t')) - p(t')$$

Truthfulness condition

$$\forall t, t' \in \mathcal{T} \;\; U(t', t) \leq U(t, t)$$

Mechanism Design

Single-player mechanism:

- Outcome space $\mathcal{O}$ *possible allocations*
- Type space $\mathcal{T} = (\mathcal{O} \to \mathbb{R})$ *valuation functions*
- Allocation rule $a : \mathcal{T} \to \mathcal{O}$ *reports to outcomes*
- Payment rule $p : \mathcal{T} \to \mathbb{R}$ *reports to payments*

Bidder with type t who reports $t' \in \mathcal{T}$ has net utility

$$U(t', t) = t(a(t')) - p(t')$$

Truthfulness condition

$$\forall t, t' \in \mathcal{T} \;\; U(t', t) \leq U(t, t)$$

Mechanism Design

Single-player mechanism:

- Outcome space $\mathcal{O}$ *possible allocations*
- Type space $\mathcal{T} = (\mathcal{O} \to \mathbb{R})$ *valuation functions*
- Allocation rule $a : \mathcal{T} \to \mathcal{O}$ *reports to outcomes*
- Payment rule $p : \mathcal{T} \to \mathbb{R}$ *reports to payments*

Bidder with type t who reports $t' \in \mathcal{T}$ has net utility

$$U(t', t) = t(a(t')) - p(t')$$

Truthfulness condition

$$\forall t, t' \in \mathcal{T} \;\; U(t', t) \leq U(t, t)$$

Myerson 1981

For single-parameter mechanisms:

Theorem

α is implementable $\iff$ α is monotone

Implementable means payments p making (α, p) truthful

Equivalently:

Theorem

α is implementable $\iff$ $\exists G : \mathcal{T} \to \mathbb{R}$ convex s.t. α is a subgradient to G

G is the consumer surplus

Myerson 1981

For single-parameter mechanisms:

Theorem

α is implementable $\iff$ α is monotone

Implementable means payments p making (α, p) truthful

Equivalently:

Theorem

α is implementable $\iff$ $\exists G : \mathcal{T} \to \mathbb{R}$ convex s.t. α is a subgradient to G

G is the consumer surplus

Scoring Rules

- Outcome space $\mathcal{O}$ *mutually exclusive events*
- Private belief $p \in \Delta_{\mathcal{O}}$ *probabilities over outcomes*
- Scoring rule $S : \Delta_{\mathcal{O}} \times \mathcal{O} \to \mathbb{R}$ *score of report given an outcome*

Expected score of report p' given truth p is

$$S(p', p) := \mathop{\mathbb{E}}_{o \sim p} [S(p', o)]$$

Truthfulness condition

$$\forall p, p' \;\; S(p', p) \leq S(p, p)$$

Scoring Rules

- Outcome space $\mathcal{O}$ *mutually exclusive events*
- Private belief $p \in \Delta_{\mathcal{O}}$ *probabilities over outcomes*
- Scoring rule $S : \Delta_{\mathcal{O}} \times \mathcal{O} \to \mathbb{R}$ *score of report given an outcome*

Expected score of report p' given truth p is

$$S(p', p) := \mathbb{E}_{o \sim p}\left[S(p', o)\right]$$

Truthfulness condition

$$\forall p, p' \quad S(p', p) \leq S(p, p)$$

Scoring Rules

- Outcome space $\mathcal{O}$ *mutually exclusive events*
- Private belief $p \in \Delta_{\mathcal{O}}$ *probabilities over outcomes*
- Scoring rule $S : \Delta_{\mathcal{O}} \times \mathcal{O} \to \mathbb{R}$ *score of report given an outcome*

Expected score of report p' given truth p is

$$S(p', p) := \mathop{\mathbb{E}}_{o \sim p} [S(p', o)]$$

Truthfulness condition

$$\forall p, p' \;\; S(p', p) \leq S(p, p)$$

Gneiting and Raftery 2007

Theorem

Scoring rule S is truthful $\iff$ there is some convex $G : \Delta_{\mathcal{O}} \to \mathbb{R}$ with subgradients $\{dG_p\}$ such that

$$S(p, o) = G(p) + dG_p(\mathbf{1}_o - p)$$

What's the Connection?

Mechanism:

- Outcomes $\mathcal{O}$
- Type $\mathcal{T} = (\mathcal{O} \to \mathbb{R})$
- Utility $U(t', t)$

Scoring rule:

- Outcomes $\mathcal{O}$
- Belief $p \in \Delta_{\mathcal{O}}$
- Score $S(p', p)$

Truthfulness

$$U(t', t) \leq U(t, t)$$

Truthfulness

$$S(p', p) \leq S(p, p)$$

$$U(t', t) = t(a(t')) - p(t') = \langle t, \mathbf{1}_{a(t')} \rangle - p(t')$$

$$S(p', p) := \mathop{\mathbb{E}}_{o \sim p} [S(p', o)] = \langle p, S(p', \cdot) \rangle$$

Reward *affine* in private info!

What's the Connection?

Mechanism:

- Outcomes $\mathcal{O}$
- Type $\mathcal{T} = (\mathcal{O} \to \mathbb{R})$
- Utility $U(t', t)$

Scoring rule:

- Outcomes $\mathcal{O}$
- Belief $p \in \Delta_{\mathcal{O}}$
- Score $S(p', p)$

Truthfulness

$$U(t', t) \le U(t, t)$$

Truthfulness

$$S(p', p) \le S(p, p)$$

$$U(t', t) = t(a(t')) - p(t')$$
$$= \langle t, \mathbb{1}_{a(t')} \rangle - p(t')$$

$$S(p', p) := \mathop{\mathbb{E}}_{o \sim p} [S(p', o)]$$
$$= \langle p, S(p', \cdot) \rangle$$

Reward *affine* in private info!

What's the Connection?

Mechanism:

- Outcomes $\mathcal{O}$
- Type $\mathcal{T} = (\mathcal{O} \to \mathbb{R})$
- Utility $U(t', t)$

Truthfulness

$$U(t', t) \leq U(t, t)$$

Scoring rule:

- Outcomes $\mathcal{O}$
- Belief $p \in \Delta_{\mathcal{O}}$
- Score $S(p', p)$

Truthfulness

$$S(p', p) \leq S(p, p)$$

$$U(t', t) = t(a(t')) - p(t')$$

$$= \langle t, \mathbb{1}_{a(t')} \rangle - p(t')$$

$$S(p', p) := \mathop{\mathbb{E}}_{o \sim p} [S(p', o)]$$

$$= \langle p, S(p', \cdot) \rangle$$

Reward *affine* in private info!

What's the Connection?

Mechanism:

- Outcomes $\mathcal{O}$
- Type $\mathcal{T} = (\mathcal{O} \to \mathbb{R})$
- Utility $U(t', t)$

Scoring rule:

- Outcomes $\mathcal{O}$
- Belief $p \in \Delta_{\mathcal{O}}$
- Score $S(p', p)$

Truthfulness

$$U(t', t) \le U(t, t)$$

Truthfulness

$$S(p', p) \le S(p, p)$$

$$U(t', t) = t(a(t')) - p(t')$$

$$= \langle t, \mathbf{1}_{a(t')} \rangle - p(t')$$

$$S(p', p) := \mathop{\mathbb{E}}_{o \sim p} [S(p', o)]$$

$$= \langle p, S(p', \cdot) \rangle$$

Reward *affine* in private info!

What's the Connection?

Mechanism:

- Outcomes $\mathcal{O}$
- Type $\mathcal{T} = (\mathcal{O} \to \mathbb{R})$
- Utility $U(t', t)$

Scoring rule:

- Outcomes $\mathcal{O}$
- Belief $p \in \Delta_{\mathcal{O}}$
- Score $S(p', p)$

Truthfulness
$U(t', t) \leq U(t, t)$

Truthfulness
$S(p', p) \leq S(p, p)$

$$U(t', t) = t(a(t')) - p(t') = \langle t, \mathbf{1}_{a(t')} \rangle - p(t')$$

$$S(p', p) := \mathop{\mathbb{E}}_{o \sim p} [S(p', o)] = \langle p, S(p', \cdot) \rangle$$

Reward *affine* in private info!

Our Model: Affine Score

- Type space $\mathcal{T}$ *any subset of a vector space*
- Reward space $\mathcal{A} \subseteq \text{Aff}(\mathcal{T} \to \mathbb{R})$ *affine functions on types*
- Affine score $S : \mathcal{T} \to \mathcal{A}$

Truthfulness condition

$$S(t')(t) \leq S(t)(t)$$

Observation: $G(t) := \sup_{t'} S(t')(t)$ convex

and S truthful $\implies G(t) = S(t)(t)$

Our Model: Affine Score

- Type space $\mathcal{T}$ *any subset of a vector space*
- Reward space $\mathcal{A} \subseteq \text{Aff}(\mathcal{T} \to \mathbb{R})$ *affine functions on types*
- Affine score $S : \mathcal{T} \to \mathcal{A}$

Truthfulness condition

$$S(t')(t) \leq S(t)(t)$$

Observation: $G(t) := \sup_{t'} S(t')(t)$ convex

and S truthful $\implies G(t) = S(t)(t)$

Our Model: Affine Score

- Type space $\mathcal{T}$ *any subset of a vector space*
- Reward space $\mathcal{A} \subseteq \mathrm{Aff}(\mathcal{T} \to \mathbb{R})$ *affine functions on types*
- Affine score $S : \mathcal{T} \to \mathcal{A}$

Truthfulness condition

$$S(t')(t) \leq S(t)(t)$$

Observation: $G(t) := \sup_{t'} S(t')(t)$ convex

and S truthful $\implies G(t) = S(t)(t)$

A General Truthfulness Characterization

Theorem

Affine score $S : \mathcal{T} \to \mathcal{A}$ is truthful if and only if there exists some convex $G : Conv(\mathcal{T}) \to \mathbb{R}$, and subgradients $\{dG_t\}$, such that

$$S(t')(t) = G(t') + dG_{t'}(t - t').$$

- Techniques from Gneiting-Raftery and Archer-Kleinberg
- Immediately gives previous scoring rule and mechanism characterizations

Proof: Convex $\mathcal{T}$

Proof of $\Longleftarrow$:

- $S(t')(t) = G(t') + dG_{t'}(t - t')$
  $\leq G(t) = S(t,t)$ *by def. subgradient*

Proof of $\Longrightarrow$:

- $G(t) := \sup_{t'} S(t')(t)$ *convex as pointwise supremum*
- Define $dG_t(\cdot) = S_\ell(t)(\cdot)$ *linear part of $S(t)$*
- $S(t)(\cdot)$ subgradient to G at t: *by truthfulness*

$$G(t') + dG_{t'}(t - t') = S(t')(t) \leq G(t)$$

Proof: Convex $\mathcal{T}$

Proof of $\Longleftarrow$:

- $S(t')(t) = G(t') + dG_{t'}(t - t')$
  $\leq G(t) = S(t, t)$ *by def. subgradient*

Proof of $\Longrightarrow$:

- $G(t) := \sup_{t'} S(t')(t)$ *convex as pointwise supremum!*
- Define $dG_t(\cdot) = S_\ell(t)(\cdot)$ *linear part of $S(t)$*
- $S(t)(\cdot)$ subgradient to G at t: *by truthfulness*

$$G(t') + dG_{t'}(t - t') \; = \; S(t')(t) \; \leq \; G(t)$$

Proof: Convex $\mathcal{T}$

Proof of $\Longleftarrow$:

- $S(t')(t) = G(t') + dG_{t'}(t - t')$
  $\leq G(t) = S(t, t)$ *by def. subgradient*

Proof of $\Longrightarrow$:

- $G(t) := \sup_{t'} S(t')(t)$ *convex as pointwise supremum!*
- Define $dG_t(\cdot) = S_\ell(t)(\cdot)$ *linear part of $S(t)$*
- $S(t)(\cdot)$ subgradient to G at t: *by truthfulness*

$$G(t') + dG_{t'}(t - t') = S(t')(t) \leq G(t)$$

Proof: Convex $\mathcal{T}$

Proof of $\Longleftarrow$:

- $S(t')(t) = G(t') + dG_{t'}(t - t')$
  $\leq G(t) = S(t,t)$ *by def. subgradient*

Proof of $\Longrightarrow$:

- $G(t) := \sup_{t'} S(t')(t)$ *convex as pointwise supremum!*
- Define $dG_t(\cdot) = S_\ell(t)(\cdot)$ *linear part of $S(t)$*
- $S(t)(\cdot)$ subgradient to G at t: *by truthfulness*

$$G(t') + dG_{t'}(t - t') \;=\; S(t')(t) \;\leq\; G(t)$$

Proof: Non-Convex $\mathcal{T}$

Proof of $\Longleftarrow$: same.

Proof of $\Longrightarrow$:

- Consider $\hat{t} \in Conv(\mathcal{T}) \setminus \mathcal{T}$
- Write $\hat{t} = \sum_i \alpha_i t_i$ for $t_i \in \mathcal{T}$
- Define $S(t)(\hat{t}) = \sum_i \alpha_i S(t)(t_i)$
- Define $G(\hat{t}) = \sup_{t \in \mathcal{T}} S(t)(\hat{t})$

Proceed as before...

Proof: Non-Convex $\mathcal{T}$

Proof of $\Longleftarrow$: same.

Proof of $\Longrightarrow$:

- Consider $\hat{t} \in Conv(\mathcal{T}) \setminus \mathcal{T}$
- Write $\hat{t} = \sum_i \alpha_i t_i$ for $t_i \in \mathcal{T}$
- Define $S(t)(\hat{t}) = \sum_i \alpha_i S(t)(t_i)$
- Define $G(\hat{t}) = \sup_{t \in \mathcal{T}} S(t)(\hat{t})$

Proceed as before...

Proof: Non-Convex $\mathcal{T}$

Proof of $\Longleftarrow$: same.

Proof of $\Longrightarrow$:

- Consider $\hat{t} \in Conv(\mathcal{T}) \setminus \mathcal{T}$
- Write $\hat{t} = \sum_i \alpha_i t_i$ for $t_i \in \mathcal{T}$
- Define $S(t)(\hat{t}) = \sum_i \alpha_i S(t)(t_i)$
- Define $G(\hat{t}) = \sup_{t \in \mathcal{T}} S(t)(\hat{t})$

Proceed as before...

Proof: Non-Convex $\mathcal{T}$

Proof of $\Longleftarrow$: same.

Proof of $\Longrightarrow$:

- Consider $\hat{t} \in Conv(\mathcal{T}) \setminus \mathcal{T}$
- Write $\hat{t} = \sum_i \alpha_i t_i$ for $t_i \in \mathcal{T}$
- Define $S(t)(\hat{t}) = \sum_i \alpha_i S(t)(t_i)$
- Define $G(\hat{t}) = \sup_{t \in \mathcal{T}} S(t)(\hat{t})$

Proceed as before...

Proof: Non-Convex $\mathcal{T}$

Proof of $\Longleftarrow$: same.

Proof of $\Longrightarrow$:

- Consider $\hat{t} \in Conv(\mathcal{T}) \setminus \mathcal{T}$
- Write $\hat{t} = \sum_i \alpha_i t_i$ for $t_i \in \mathcal{T}$
- Define $S(t)(\hat{t}) = \sum_i \alpha_i S(t)(t_i)$
- Define $G(\hat{t}) = \sup_{t \in \mathcal{T}} S(t)(\hat{t})$

Proceed as before...

Proof: Non-Convex $\mathcal{T}$

Proof of $\Longleftarrow$: same.

Proof of $\Longrightarrow$:

- Consider $\hat{t} \in Conv(\mathcal{T}) \setminus \mathcal{T}$
- Write $\hat{t} = \sum_i \alpha_i t_i$ for $t_i \in \mathcal{T}$
- Define $S(t)(\hat{t}) = \sum_i \alpha_i S(t)(t_i)$
- Define $G(\hat{t}) = \sup_{t \in \mathcal{T}} S(t)(\hat{t})$

Proceed as before...

Immediate New Results

1. Proper scoring rules for non-convex sets of distributions
   Fewer constraints $\Longrightarrow$ more scoring rules?

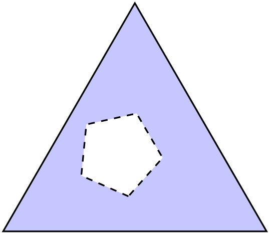

2. “Local” mechanisms and scoring rules
   Convexity is a local property

Immediate New Results

1. Proper scoring rules for non-convex sets of distributions

   Fewer constraints $\implies$ more scoring rules?

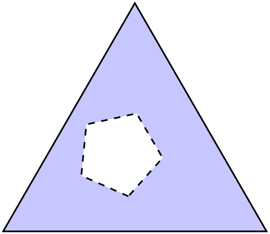

2. "Local" mechanisms and scoring rules

   Convexity is a local property

Mechanism Design: Implementability of α

Definition

$\{dG_t\}_{t\in\mathcal{T}}$ satisfies *cyclic monotonicity (CMON)* if for all finite sets $\{t_0,\ldots,t_k\}\subseteq\mathcal{T}$,

$$\sum_{i=0}^{k} dG_{t_i}(t_{i+1}-t_i)\le 0.$$

CMON with $k=2$ is *Weak monotonicity (WMON)*.

Let $L_{xy}=\int_0^1 dG_{\beta y+(1-\beta)x}\,(y-x)\,d\beta$.

Definition

$\{dG_t\}_{t\in\mathcal{T}}$ satisfies *path independence (PI)* if for all $x,y,z\in\mathcal{T}$

$$L_{xy}+L_{yz}=L_{xz}$$

Previous Characterizations

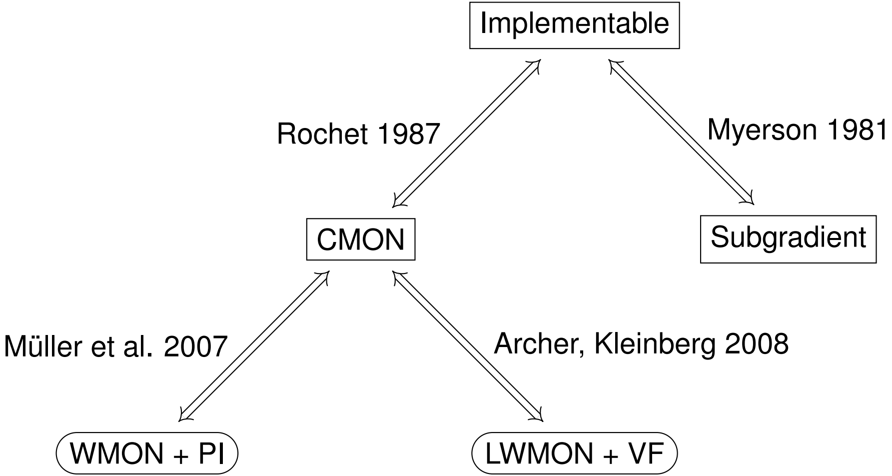

A New Proof Structure

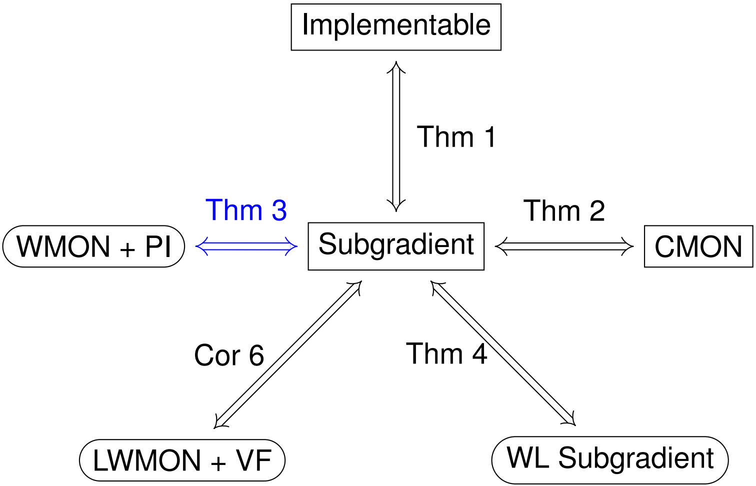

Reproving Müller et al.

WMON + PI $\xrightarrow{\text{Thm 3}}$ Subgradient

New proof via construction of G:

- Fix $G(t_0)$
- Extend $G(t) = L_{t_0 t}$ *integrable by WMON, consistent by PI*
- Subgradient by simple computation

Reproving Müller et al.

WMON + PI $\xrightarrow{\text{Thm 3}}$ Subgradient

New proof via construction of G:

- Fix $G(t_0)$
- Extend $G(t) = L_{t_0\,t}$ *integrable by WMON, consistent by PI*
- Subgradient by simple computation

Reproving Müller et al.

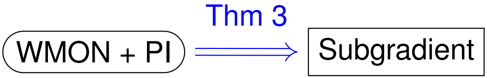

New proof via construction of G:

- Fix $G(t_0)$
- Extend $G(t) = L_{t_0\,t}$ *integrable by WMON, consistent by PI*
- Subgradient by simple computation

TMI

Q: What if types are exponential (or infinite!) in size?

A: Use summary information / low-dim representation

Examples:

- Scoring rules for statistics *[Lambert-Pennock-Shoham, Gneiting]*
- Rankings instead of utilities *[Carroll]*
- ...

TMI

Q: What if types are exponential (or infinite!) in size?

A: Use summary information / low-dim representation

Examples:

- Scoring rules for statistics *[Lambert-Pennock-Shoham, Gneiting]*
- Rankings instead of utilities *[Carroll]*
- ...

TMI

Q: What if types are exponential (or infinite!) in size?

A: Use summary information / low-dim representation

Examples:

- Scoring rules for statistics *[Lambert-Pennock-Shoham, Gneiting]*
- Rankings instead of utilities *[Carroll]*
- ...

More Formally...

Wish to change report space from $\mathcal{T}$ to some other R

$$S : R \to \mathrm{Aff}(\mathcal{T} \to \mathbb{R}); \qquad S(r)(t)$$

What does truthful mean now?

Properties

Definition

A *property* is a map $\Gamma : \mathcal{T} \to R$ specifying the correct report $r = \Gamma(t)$ for each type t.

Truthfulness condition

$$S(r')(t) \leq S(\Gamma(t))(t)$$

We say S *elicits* Γ.

Properties

Definition

A *property* is a map $\Gamma : \mathcal{T} \rightarrow R$ specifying the correct report $r = \Gamma(t)$ for each type t.

Truthfulness condition

$$S(r')(t) \leq S(\Gamma(t))(t)$$

We say S *elicits* Γ.

A New Result

Theorem

Property Γ is elicitable iff there exists $G : \mathcal{T} \to \mathbb{R}$ differentiable and convex, and map $\varphi : R \to \nabla G(\mathcal{T})$, such that $\varphi(\Gamma(t)) = \nabla G(t)$.

New insights:

- Elicitable properties == subgradients!
- Properties specify where G should be *flat*

A New Result

Theorem

Property Γ is elicitable iff there exists $G : \mathcal{T} \to \mathbb{R}$ differentiable and convex, and map $\varphi : R \to \nabla G(\mathcal{T})$, such that $\varphi(\Gamma(t)) = \nabla G(t)$.

New insights:

- Elicitable properties == subgradients!
- Properties specify where G should be *flat*

Finite R: Power Diagram

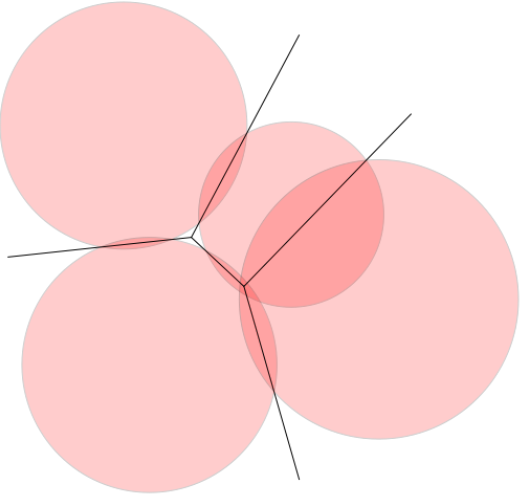

Cells = types with same report. Application: rankings!

Thanks!