Reducing (Bayesian) Mechanism Design to Algorithm Design

Bobby Kleinberg
CWI Workshop on Advances in Algorithmic Game Theory
September 3, 2010

Joint work with: Jason Hartline, Azarakhsh Malekian

FLEXFLIX

Petra’s Queue

1	Into Great Silence	High
2	Lulu & Jimi	Med.
3	Avatar	Low
4	The White Ribbon	High
5	Signs of Life	Low

FLEXFLIX

Bobby’s Queue

1	Transformers 2	Med.
2	GI Joe: Rise of Cobra	Med.
3	The Secret in Their Eyes	Low
4	Avatar	High
5	Sherlock Holmes	High

Petra's Queue

Into Great Silence	High
Lulu & Jimi	Med.
Avatar	Low
The White Ribbon	High
Signs of Life	Low

Bobby's Queue

Transformers 2	Med.
GI Joe: Rise of Cobra	Med.
The Secret in Their Eyes	Low
Avatar	High
Sherlock Holmes	High

Petra's Queue

Into Great Silence	High
Lulu & Jimi	Med.
Avatar	Low
The White Ribbon	High
Signs of Life	Low

Bobby's Queue

Transformers 2	Med.
GI Joe: Rise of Cobra	Med.
The Secret in Their Eyes	Low
Avatar	High
Sherlock Holmes	High

Petra's Queue

Into Great Silence	High
Lulu & Jimi	Med.
The White Ribbon	High
Avatar	Low
Signs of Life	Low

Bobby's Queue

Transformers 2	Med.
GI Joe: Rise of Cobra	Med.
Avatar	High
The Secret in Their Eyes	Low
Sherlock Holmes	High

Petra's Queue

Bobby's Queue

Petra's Queue		Bobby's Queue	
Into Great Silence	High	Transformers 2	Med.
Lulu & Jimi	Med.	GI Joe: Rise of Cobra	Med.
The White Ribbon	High	Avatar	High
Avatar	Low	The Secret in Their Eyes	Low
Signs of Life	Low	Sherlock Holmes	High

Petra's Queue

Bobby's Queue

The White Ribbon	High
Avatar	Low
Speed Racer	Low
Signs of Life	High
Transformers 2	Low

Avatar	High
The Secret in Their Eyes	Low
Sherlock Holmes	High
Iron Man 2	High
Inception	High

Petra’s Queue

The White Ribbon	High
Avatar	Low
Speed Racer	Low
Signs of Life	High
Transformers 2	Low

Bobby’s Queue

Avatar	High
The Secret in Their Eyes	Low
Sherlock Holmes	High
Iron Man 2	High
Inception	High

Our Guiding Question

Is there a reduction that makes an arbitrary algorithm incentive compatible, with little or no loss in social welfare?

Preliminaries

X_i = type space $\qquad v_i : X_i \times Y \to \mathbf{R}$ (valuation)

$X = X_1 \times \cdots \times X_n$ $\qquad f : X \to Y$ (allocation)

Y = outcomes $\qquad p_i : X \to \mathbf{R}$ (payment)

Objective: $\max_y \{ \sum v_i(x_i, y) \}$ "social welfare"

Truthfulness: Two different notions...

Preliminaries

X_i = type space $\quad v_i : X_i \times Y \to \mathbf{R}$ (valuation)

$X = X_1 \times \cdots \times X_n$ $\quad f : X \to Y$ (allocation)

Y = outcomes $\quad p_i : X \to \mathbf{R}$ (payment)

Objective: $\max_y \{ \sum v_i(x_i,y) \}$ "social welfare"

Truthfulness: Two different notions...

Dominant Strategy $\forall i\ \forall x_{-i}\ \ v_i(x_i,f(\cdot,x_{-i}))-p_i(\cdot)$ max@x_i

Bayesian $\forall i\ \ \mathbf{E}[v_i(x_i,f(\cdot,x_{-i}))]-p_i(\cdot)$ max@x_i

Cyclic Monotonicity

x_1 $y_1=f(x_1)$

x_2 $y_2=f(x_2)$

x_3 $y_3=f(x_3)$

x_4 $y_4=f(x_4)$

x_5 $y_5=f(x_5)$

Truthfulness of single-player mechanism (f,p) implies “max matching property” of f.

Converse: $\exists$p making (f,p) truthful if f satisfies the max matching property, a.k.a. cyclic monotonicity (CMON).

Mechanism design is algorithm design with a cyclic monotonicity constraint.

Cyclic Monotonicity

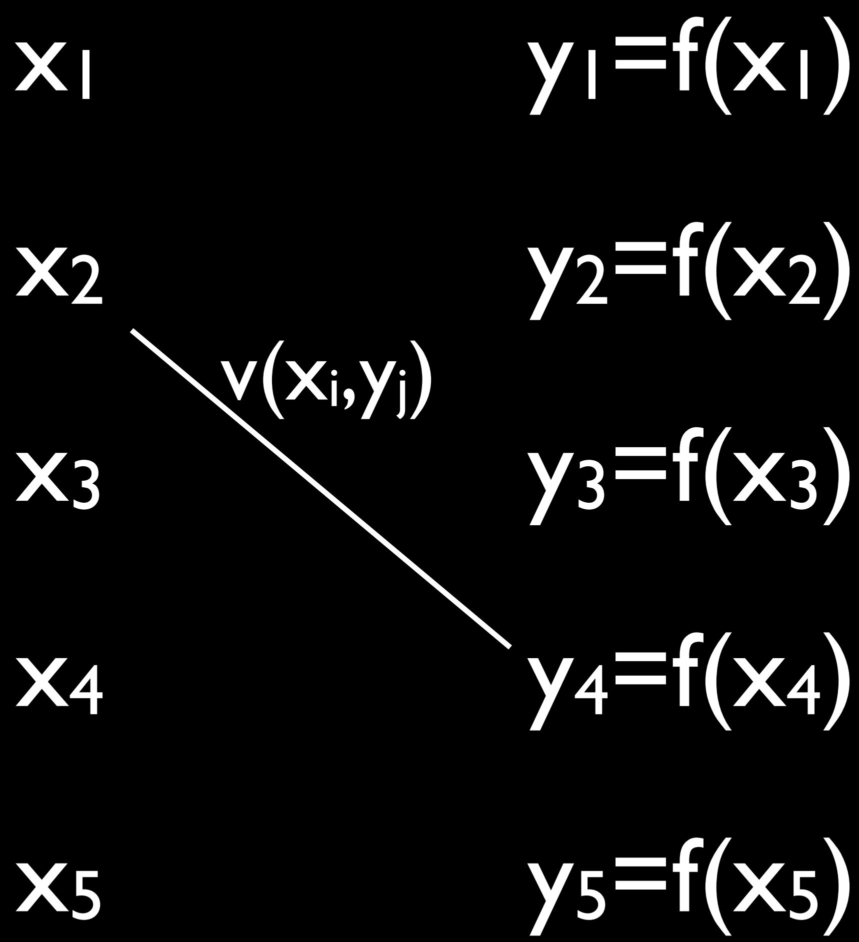

Truthfulness of single-player mechanism (f,p) implies “max matching property” of f.

Converse: ∃p making (f,p) truthful if f satisfies the max matching property, a.k.a. cyclic monotonicity (CMON).

Mechanism design is algorithm design with a cyclic monotonicity constraint.

Cyclic Monotonicity

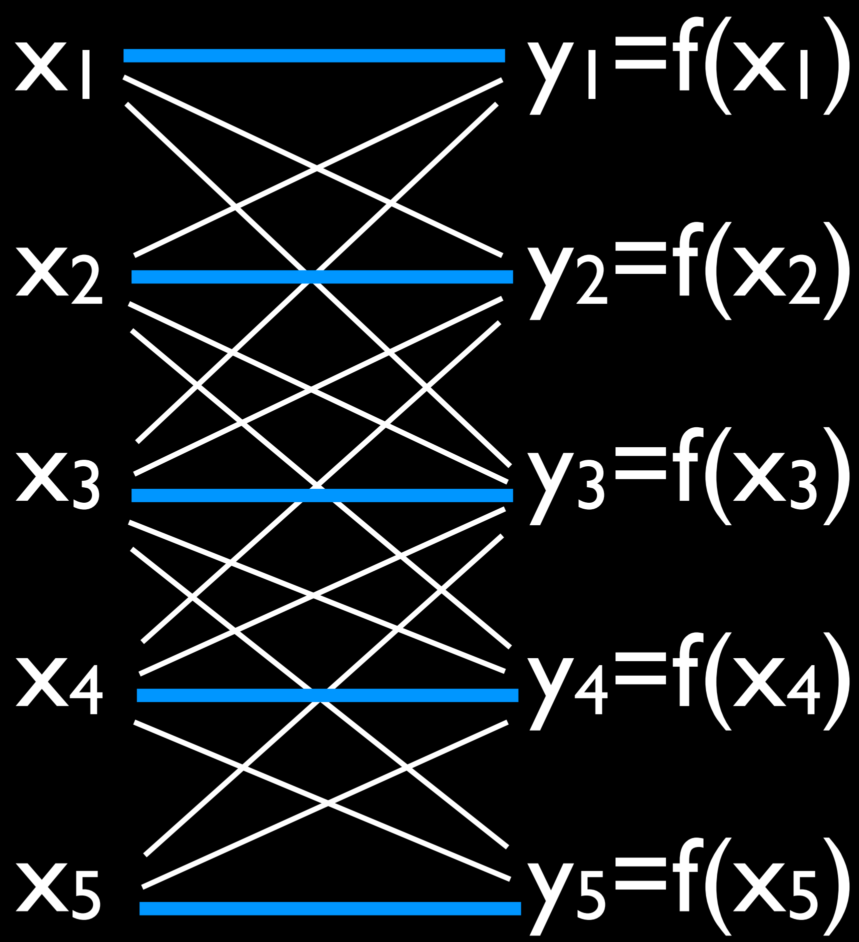

Truthfulness of single-player mechanism (f,p) implies “max matching property” of f.

Converse: ∃p making (f,p) truthful if f satisfies the max matching property, a.k.a. cyclic monotonicity (CMON).

Mechanism design is algorithm design with a cyclic monotonicity constraint.

Our Guiding Question

Is there a reduction that makes an arbitrary algorithm incentive compatible, with little or no loss in social welfare?

Our Guiding Question

Is there a reduction that makes an arbitrary algorithm incentive compatible, with little or no loss in social welfare?

YES, VCG.

Our Guiding Question

*Is there an **efficient** reduction that makes an arbitrary algorithm incentive compatible, with little or no loss in social welfare?*

Our Guiding Question

Is there an ***efficient*** *reduction that makes an arbitrary algorithm incentive compatible, with little or no loss in social welfare?*

Dominant Strategy **NO**

[work of Dobzinski, Lavi, Mu'alem, Nisan, Papadimitriou, Schapira, Singer, many others]

Bayesian **YES!**

[Hartline-Lucier STOC'10, this talk]

Assumptions

How are players' bid distributions specified?

1. General oracles for sampling x_i, evaluating v_i
2. Single-parameter $X_i \subseteq \mathbf{R}$, $Y \subseteq \mathbf{R}^n$, $v_i(x_i,y) = x_i \cdot y_i$
3. Discrete sample space Ω input size = $|\Omega|$

How is the algorithm f specified?

1. Black box model oracle for evaluating f
2. Ideal model additional oracle for $\mathbf{E}[f(x,x_{-i})]$

Summary of Results

Assume $v_i : X_i \times Y \to [0,1]$, and seek ε-additive approximation to social welfare of f.

	Discrete	1-Param.	General
Ideal	$\lvert\Omega\rvert$	$O(\varepsilon^{-1})$	$O(\varepsilon^{-\Delta-2})^{*}$
Black Box	$\tilde{O}(n^3\lvert\Omega\rvert^7\varepsilon^{-3})$	$\tilde{O}(\varepsilon^{-3})$	$\tilde{O}(\varepsilon^{-3\Delta-6})^{**}$

Table gives sample complexity s.

Running time is $O(ns^3)$.

The Fine Print

Assume $v_i : X_i \times Y \to [0,1]$, and seek ε-additive approximation to social welfare of f.

	Discrete	1-Param.	General
Ideal	$\lvert\Omega\rvert$	$O(\varepsilon^{-1})$	$O(\varepsilon^{-\Delta-2})^{*}$
Black Box	$\tilde{O}(n^3\lvert\Omega\rvert^7\varepsilon^{-3})$	$\tilde{O}(\varepsilon^{-3})$	$\tilde{O}(\varepsilon^{-3\Delta-6})^{**}$

* Δ = "number of parameters to specify a type or outcome"

** Mechanism is only ε-truthful, not truthful.

The Fine Print

Assume $v_i : X_i \times Y \to [0,1]$, and seek ε-additive approximation to social welfare of f.

	Discrete	1-Param.	General
			$O(\varepsilon^{-\Delta-2})$*
			$\tilde{O}(\varepsilon^{-3\Delta-6})$**

∀ sufficiently large k, each X_i can be covered by $O(k^{\Delta})$ sets of diameter 1/k in the L_∞ metric. (Distance between types is max. difference of values they assign to the same outcome.)

* Δ = “number of parameters to specify a type or outcome”

** Mechanism is only ε-truthful, not truthful.

The Fine Print

Assume $v_i : X_i \times Y \to [0,1]$, and seek ε-additive approximation to social welfare of f.

	Discrete	1-Param.	General
Ideal	$\|\Omega\|$	$O(\varepsilon^{-1})$	$O(\varepsilon^{-\Delta-2})$*
Black Box	$\tilde{O}(n^3\|\Omega\|^7\varepsilon^{-3})$	$\tilde{O}(\varepsilon^{-3})$	$\tilde{O}(\varepsilon^{-3\Delta-6})$**

* Δ = "number of parameters to specify a type or outcome"

** Mechanism is only ε-truthful, not truthful.

The Fine Print

Assume $v_i : X_i \times Y \rightarrow [0,1]$, and seek ε-additive approximation to social welfare of f.

	Discrete	1-Param.	General
Ideal	$\lvert\Omega\rvert$	$O(\varepsilon^{-1})$	$O(\varepsilon^{-\Delta-2})^{*}$
Black Box	$\tilde{O}(n^3\lvert\Omega\rvert^7\varepsilon^{-3})$	$\tilde{O}(\varepsilon^{-3})$	$\tilde{O}(\varepsilon^{-3\Delta-6})^{**}$

* Δ = "number of parameters to specify a type or outcome"

** Mechanism is only ε-truthful, not truthful.

Idea #1: Surrogates

- Replace each bid x_i with a random surrogate $\sigma(x_i)$.
- Choose outcome $y = f(\sigma(x_1),\ldots,\sigma(x_n))$.

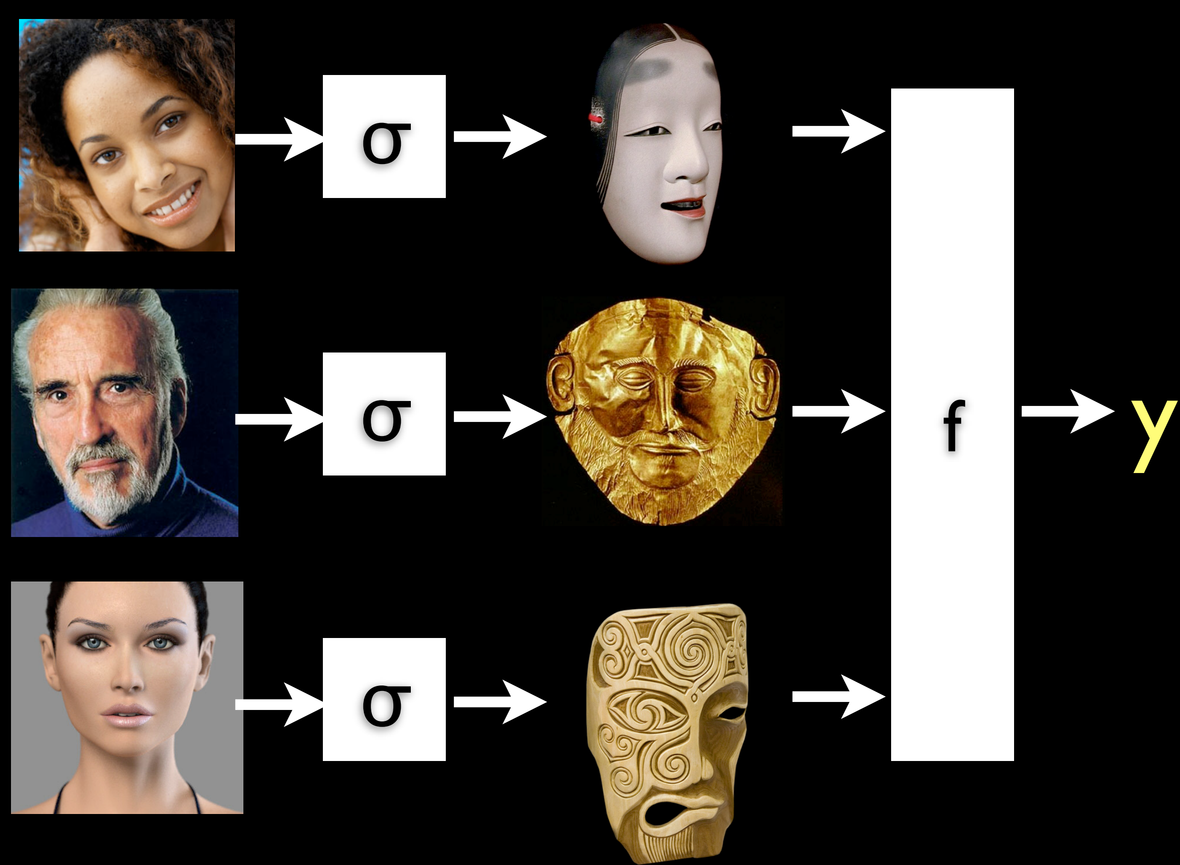

We require two properties of the sampling process $\sigma(\cdot)$.

Stationarity: stationary distrib. is the type distrib. of player i.

Monotonicity: the function $x \to \sigma(x)$ is CMON.

Idea #1: Surrogates

- Replace each bid x_i with a random surrogate $\sigma(x_i)$.
- Choose outcome $y = f(\sigma(x_1),\ldots,\sigma(x_n))$.

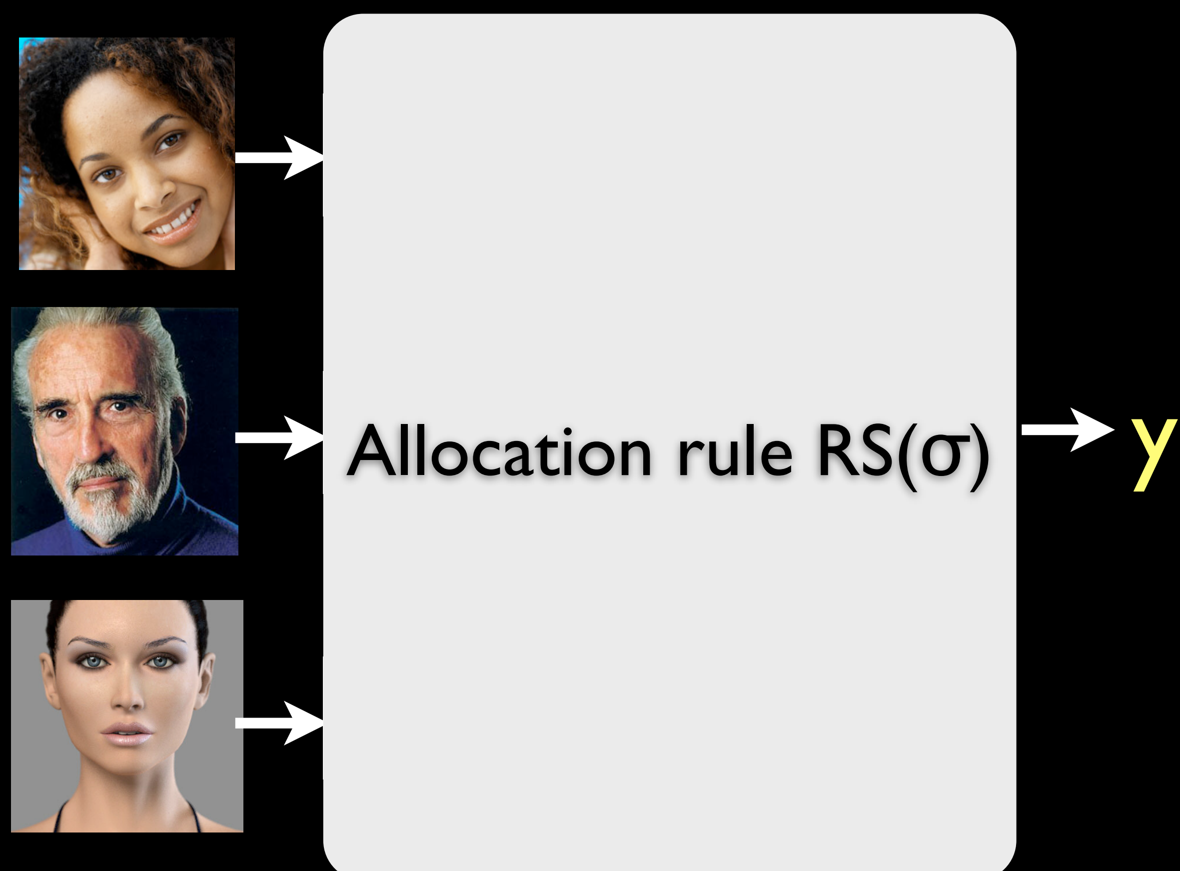

We require two properties of the sampling process $\sigma(\cdot)$.

Stationarity: stationary distrib. is the type distrib. of player i.

Monotonicity: the function $x \to \sigma(x)$ is CMON.

Idea #1: Surrogates

- Replace each bid x_i with a random surrogate $\sigma(x_i)$.
- Choose outcome $y = f(\sigma(x$

w.r.t. valuation function

$\mathcal{V}(x,y) := E[v(x, f(y, x_{-i}))]$

The expected value that type x assigns to the random outcome obtained using surrogate y.

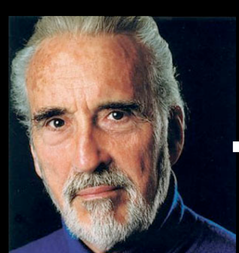

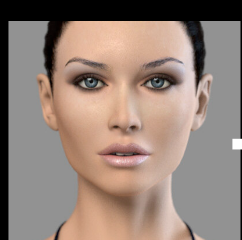

Allocation rule RS(σ) → y

is the type distrib. of player i.

Monotonicity: the function $x \to \sigma(x)$ is CMON.

Idea #1: Surrogates

Stationarity: stationary distrib. is type distrib. of player i.

Monotonicity: the function $x \to \sigma(x)$ is CMON.

Theorem: If σ satisfies these two properties, then the allocation rule $RS(\sigma)$ is CMON.

Proof:

Idea #1: Surrogates

Stationarity: stationary distrib. is type distrib. of player i.

Monotonicity: the function $x \to \sigma(x)$ is CMON.

Theorem: If σ satisfies these two properties, then the allocation rule RS(σ) is CMON.

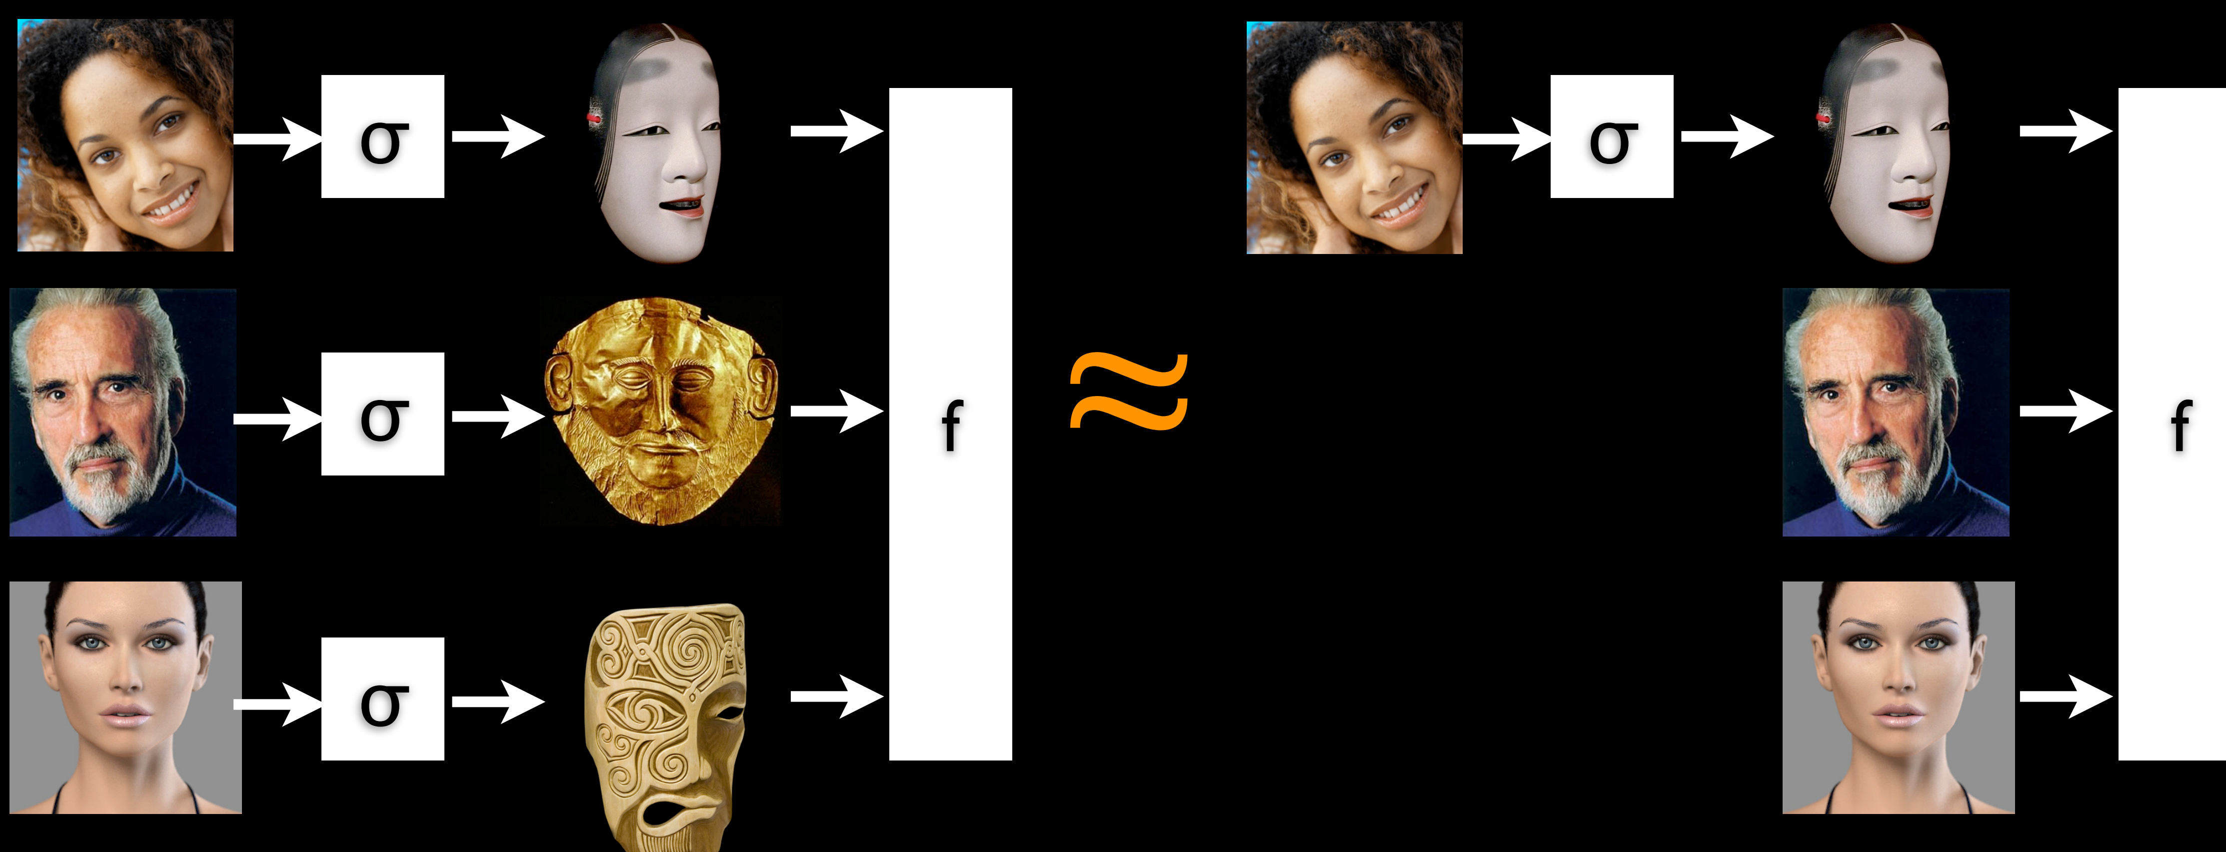

Idea #1: Surrogates

Stationarity: stationary distrib. is type distrib. of player i.

Monotonicity: the function $x \to \sigma(x)$ is CMON.

Theorem: If σ satisfies these two properties, then the allocation rule $RS(\sigma)$ is CMON.

Remark: Easy to compute payments for $RS(\sigma)$, but won't discuss the issue further in this talk.

Examples

1. σ = Id satisfies stationary, but not monotonicity unless f is monotone.
2. σ = Resample satisfies both properties, but has lousy social welfare.

Idea #2: Replicas

1. Sample *replicas* $r_1,\ldots,r_m$ and *surrogates* $s_1,\ldots,s_m$ i.i.d. from type distribution on X_i.
2. Choose random k, set $r_k = x_i$.
3. Set edge weights $w_{ij} = \mathcal{V}(r_i, s_j)$.
4. Let μ = max-weight matching.
5. Declare surrogate $\sigma(x_i) = \mu(r_k)$.






Idea #2: Replicas

1. Sample *replicas* $r_1,\ldots,r_m$ and *surrogates* $s_1,\ldots,s_m$ i.i.d. from type distribution on X_i.
2. Choose random k, set $r_k=x_i$.
3. Set edge weights $w_{ij} = \nu(r_i, s_j)$.
4. Let μ = max-weight matching.
5. Declare surrogate $\sigma(x_i) = \mu(r_k)$.






Idea #2: Replicas

1. Sample *replicas* $r_1,\ldots,r_m$ and *surrogates* $s_1,\ldots,s_m$ i.i.d. from type distribution on X_i.
2. Choose random k, set $r_k = x_i$.
3. Set edge weights $w_{ij} = \mathcal{V}(r_i, s_j)$.
4. Let μ = max-weight matching.
5. Declare surrogate $\sigma(x_i) = \mu(r_k)$.

Idea #2: Replicas

1. Sample *replicas* $r_1,\ldots,r_m$ and *surrogates* $s_1,\ldots,s_m$ i.i.d. from type distribution on X_i.
2. Choose random k, set $r_k=x_i$.
3. Set edge weights $w_{ij} = \mathcal{V}(r_i, s_j)$.
4. Let μ = max-weight matching.
5. Declare surrogate $\sigma(x_i) = \mu(r_k)$.

Idea #2: Replicas

1. Sample *replicas* $r_1,\ldots,r_m$ and *surrogates* $s_1,\ldots,s_m$ i.i.d. from type distribution on X_i.
2. Choose random k, set $r_k=x_i$.
3. Set edge weights $w_{ij} = \mathcal{V}(r_i, s_j)$.
4. Let μ = max-weight matching.
5. Declare surrogate $\sigma(x_i) = \mu(r_k)$.

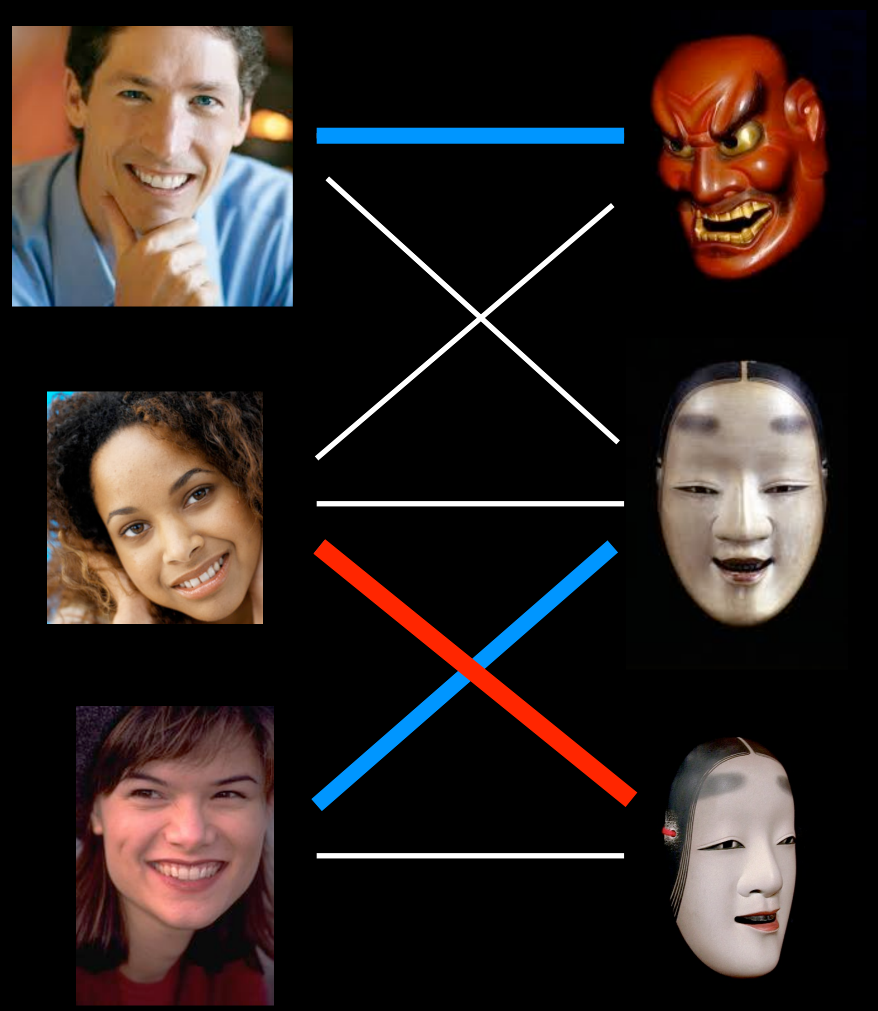

Idea #2: Replicas

1. Sample *replicas* $r_1,\ldots,r_m$ and *surrogates* $s_1,\ldots,s_m$ i.i.d. from type distribution on X_i.
2. Choose random k, set $r_k=x_i$.
3. Set edge weights $w_{ij} = \mathcal{V}(r_i, s_j)$.
4. Let μ = max-weight matching.
5. Declare surrogate $\sigma(x_i) = \mu(r_k)$.

Stationarity: Distrib. of $\mu(r_k)$ unchanged if step 2 omitted.

Idea #2: Replicas

1. Sample *replicas* $r_1,\ldots,r_m$ and *surrogates* $s_1,\ldots,s_m$ i.i.d. from type distribution on X_i.
2. Choose random k, set $r_k=x_i$.
3. Set edge weights $w_{ij} = \mathcal{V}(r_i, s_j)$.
4. Let μ = max-weight matching.
5. Declare surrogate $\sigma(x_i) = \mu(r_k)$.

Monotonicity: Conditional on replicas, surrogates, and k, the mapping from x_i to $\sigma(x_i)$ is monotone. (in fact, max'l in range)

Welfare Approximation

Welfare loss of bidder i is

$\mathcal{V}(r_k,r_k) - \mathcal{V}(r_k,\mu(r_k))$

Expectation is

$(1/m)*[\Sigma_k \mathcal{V}(r_k,r_k) - \Sigma_k \mathcal{V}(r_k,\mu(r_k))]$

This is no greater than

$(1/m)*[\Sigma_k \mathcal{V}(r_k,r_k) - \Sigma_k \mathcal{V}(r_k,\lambda(r_k))]$

for any other matching λ.

Bound this from above by

$(1/m)*[\Sigma_k ||r_k - \lambda(r_k)||_\infty]$

Choose λ to minimize the RHS.

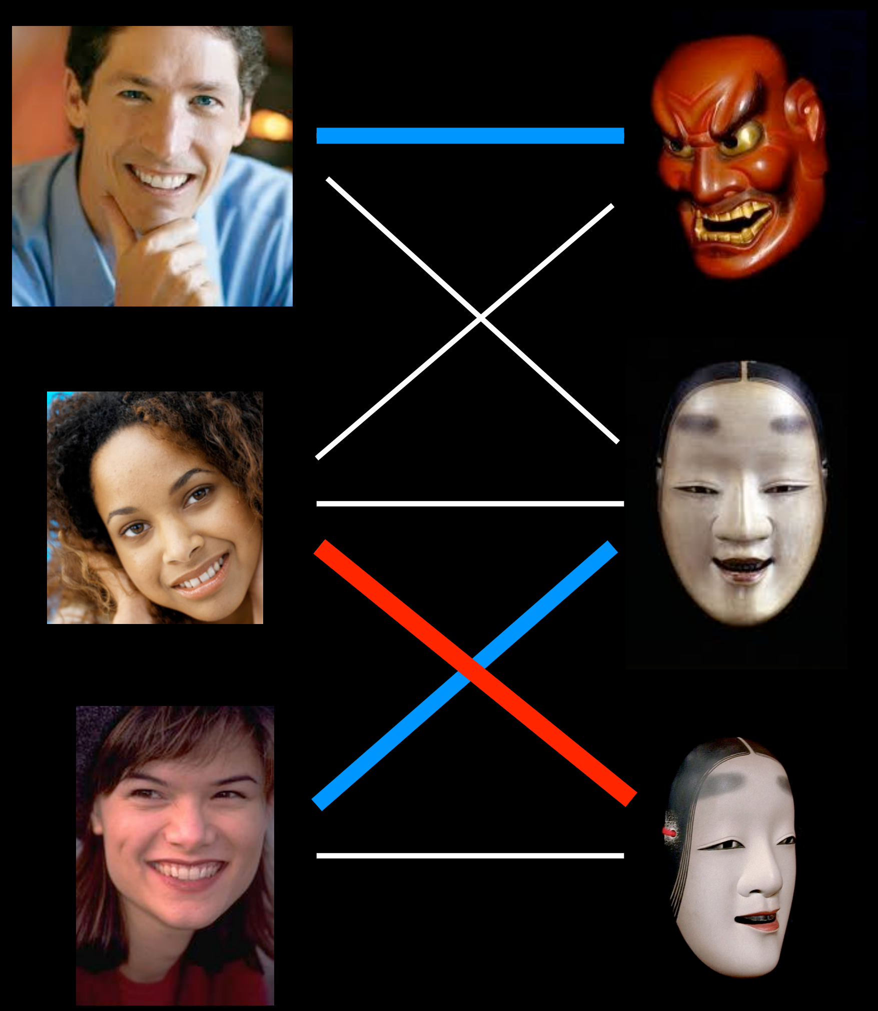

Transportation Cost

X a metric space.

Transportation cost between two m-point subsets of X is length of min-cost matching.

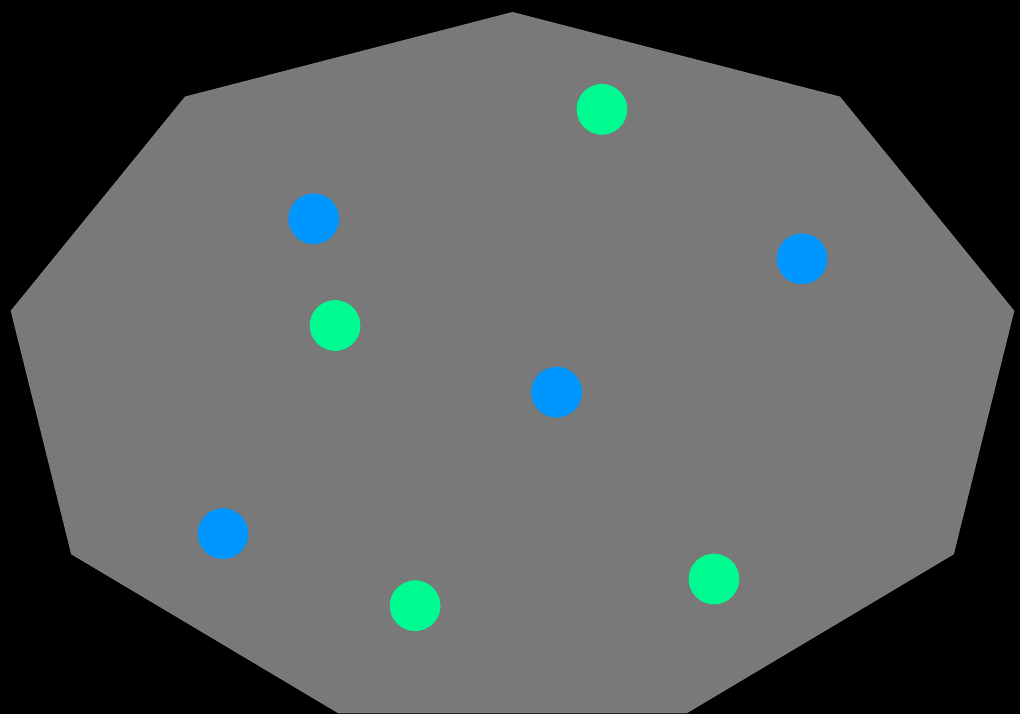

Transportation Cost

X a metric space.

Transportation cost between two m-point subsets of X is length of min-cost matching.

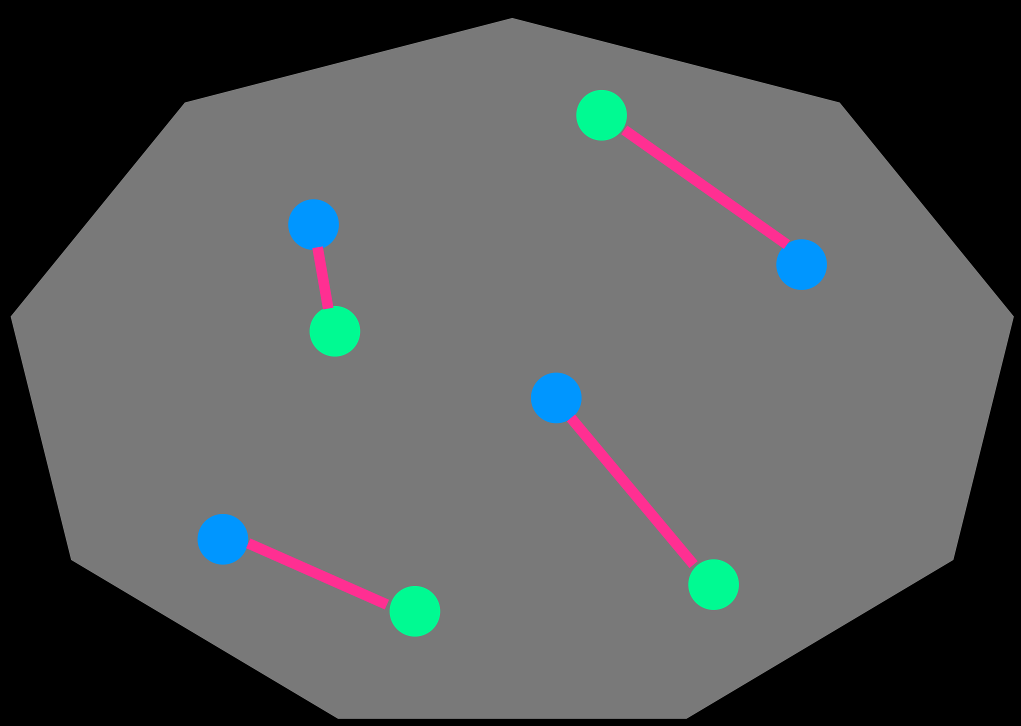

Transportation Cost

X a metric space.

Transportation cost between two m-point subsets of X is length of min-cost matching.

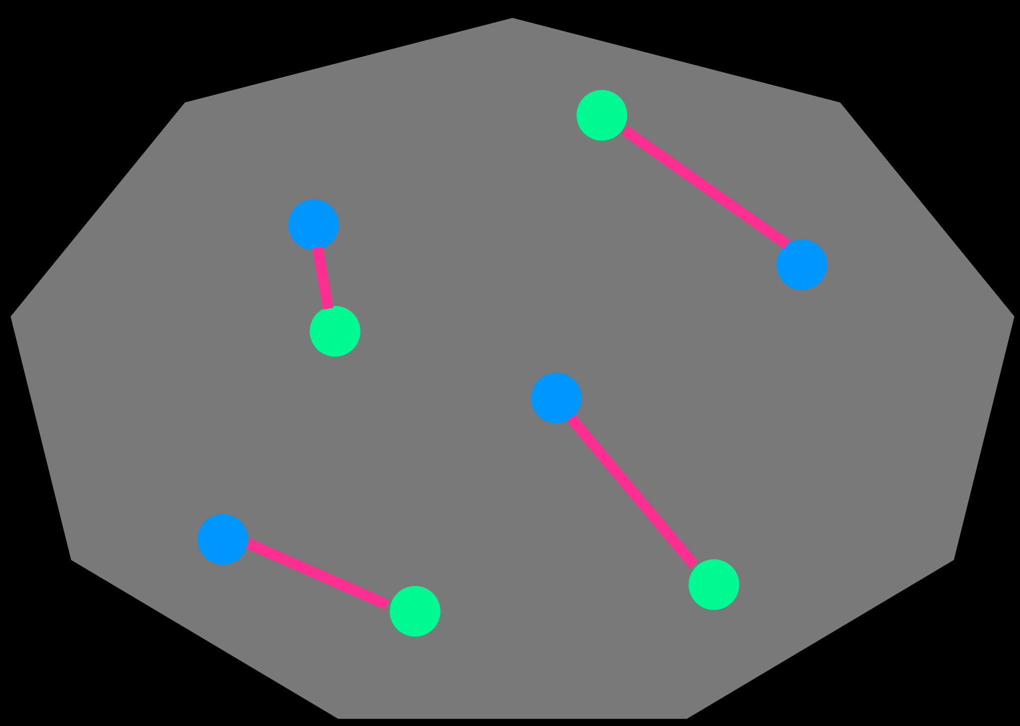

Theorem: If Diam(X)=1 and X partitions into k^{Δ} sets of diameter 1/k, the expected transportation cost of two random m-point subsets is $O(m/k + (mk^{\Delta})^{1/2})$.

Proof Sketch: Match as many points as possible to partners in same piece of partition. Bound expected number of unmatched points by $(mk^{\Delta})^{1/2}$.

Transportation Cost

X a metric space.

Transportation cost between two m-point subsets of X is length of min-cost matching.

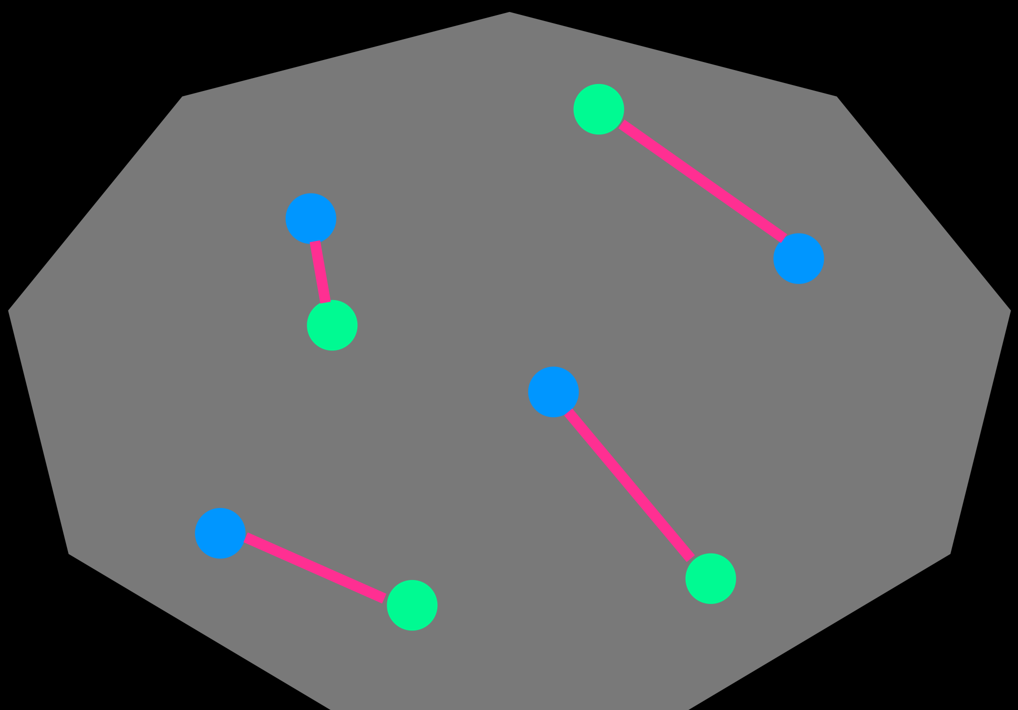

Theorem: If Diam(X)=1 and X partitions into k^{Δ} sets of diameter 1/k, the expected transportation cost of two random m-point subsets is $O(m/k + (mk^{\Delta})^{1/2})$.

Corollary: Replica-surrogate matching mechanism achieves $O(\varepsilon)$ welfare loss when $k = \varepsilon^{-1}$, $m = \varepsilon^{-\Delta-2}$.

Transportation Cost

X a metric space.

Transportation cost between two m-point subsets of X is length of min-cost matching.

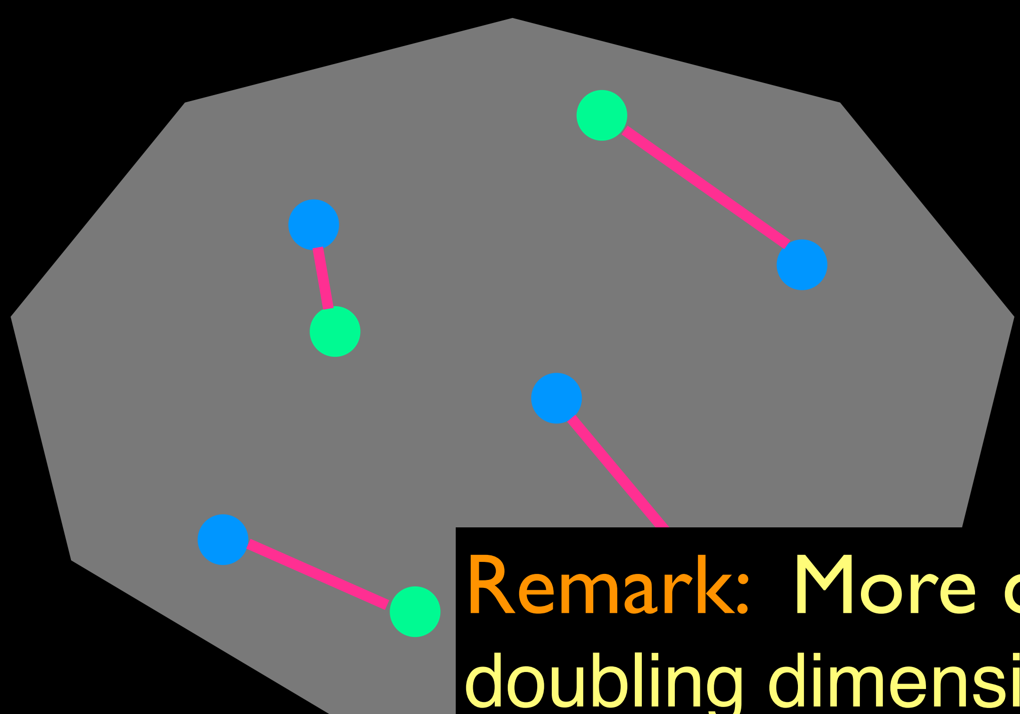

Theorem: If Diam(X)=1 and X partitions into k^Δ sets of diameter 1/k, the expected transportation cost of two random m-point subsets is $O(m/k + (mk^\Delta)^{1/2})$.

Corollary: Replica-surrogate matching mechanism achieves $O(\varepsilon)$ welfare loss when $k = \varepsilon^{-1}$, $m = \varepsilon^{-\Delta-2}$.

Remark: More careful analysis gives $m=\varepsilon^{-\Delta-1}$ in doubling dimension Δ. This is tight except for $\Delta \leq 2$.

Extensions

- Improved mechanism for single-parameter case. {Replicas} = {Surrogates}
- Mechanisms for the black box model. (Can evaluate f but can't query its exact expectation.)
  - Single-parameter case
  - Discrete type space






Extensions

- Improved mechanism for single-parameter case. {Replicas} = {Surrogates}
- Mechanisms for the black box model. (Can evaluate f but can’t query its exact expectation.)
  - Single-parameter case
  - Discrete type space






Extensions

- Improved mechanism for single-parameter case. {Replicas} = {Surrogates}
- Mechanisms for the black box model. (Can evaluate f but can't query its exact expectation.)
  - Single-parameter case
  - Discrete type space

Extensions

- Improved mechanism for single-parameter case. {Replicas} = {Surrogates}
- Mechanisms for the black box model. (Can evaluate f but can't query its exact expectation.)
  - Single-parameter case
  - Discrete type space

Extensions

- Improved mechanism for single-parameter case. {Replicas} = {Surrogates}
- Mechanisms for the black box model. (Can evaluate f but can't query its exact expectation.)
    - Single-parameter case
    - Discrete type space

Open Questions

	Discrete	1-Param.	General
Ideal	$\lvert\Omega\rvert$	$O(\varepsilon^{-1})$	$O(\varepsilon^{-\Delta-2})^{*}$
Black Box	$\tilde{O}(n^3\lvert\Omega\rvert^7\varepsilon^{-3})$	$\tilde{O}(\varepsilon^{-3})$	$\tilde{O}(\varepsilon^{-3\Delta-6})^{**}$

* Δ=covering dimension ** ε-truthful, but not truthful

- Exponential dependence on Δ necessary?
- Remove the double-asterisk ... please!!
- Achieve ε-approximation pointwise, not in expectation.
- Approximate other objectives, e.g. fairness.