Sybil loving without a trace of budget balance; not to mention core

THE CURIOUS CASE OF VCG PAYMENTS

The Curious Case of VCG Payments

	A	B	AB
X	8		
Y		8	
Z			10

The Curious Case of VCG Payments

	A	B	AB
X	8(2)		
Y		8(2)	
Z			10

The Curious Case of VCG Payments

7	6
6	3
4	2

“By how much do the incentives for truthful reporting fail when other design objectives are imposed as constraints?”

CORE SELECTING PACKAGE AUCTIONS & QUANTIFYING THE STRATEGYPROOFNESS OF MECHANISMS VIA METRICS ON PAYOFF DISTRIBUTIONS

The Quest for the (almost) Holy Grail

- VCG payments are the only strategy proof payments for package auctions or exchanges
- Unfortunately:
  - They’re not sybilproof
  - They’re not budget balanced
  - The outcome may not be in the core
- Instead, consider requisite properties and minimize incentives to misreport

Or “Dishonest games you’d like to play”

PACKAGE AUCTIONS AND CORE

Mawwiage. Mawwiage is what bwings us togetha today.

Fred
- Wilma
- Betty

Barney
- Betty
- Wilma

Wilma
- Barney
- Fred

Betty
- Fred
- Barney

Mawwiage. Mawwiage is what bwings us togetha today.

Fred
- Wilma
- Betty

Barney
- Betty
- Wilma

Wilma
- Barney
- Fred

Betty
- Fred
- Barney

Mawwiage. Mawwiage is what bwings us togetha today.

Fred
- Wilma
- Betty

Barney
- Betty
- Wilma

Wilma
- Barney
- ~~Fred~~

Betty
- Fred
- ~~Barney~~

Mawwiage. Mawwiage is what bwings us togetha today.

Fred
- Wilma
- Betty

Barney
- Betty
- Wilma

Bam Bam
- Wilma
- Pebbles (not pictured)

Wilma
- Barney
- ~~Fred~~

Betty
- Fred
- ~~Barney~~

THE FOLLOWING SLIDE COURTESY OF PROFESSOR AL ROTH

Market	Stable	Still in use (halted unraveling)
NRMP	yes	yes (new design in '98)
Edinburgh ('69)	*yes*	*yes*
Cardiff	*yes*	*yes*
Birmingham	*no*	*no*
Edinburgh ('67)	*no*	*no*
Newcastle	*no*	*no*
Sheffield	no	no
Cambridge	no	yes
London Hospital	no	yes
Medical Specialties	yes	yes (~30 markets, 1 failure)
Canadian Lawyers	yes	yes (Alberta, no BC, Ontario)
Dental Residencies	yes	yes (5) (no 2)
Osteopaths (< '94)	no	no
Osteopaths (≥ '94)	yes	yes
Pharmacists	yes	yes
Reform rabbis	yes (first used in '97-98)	yes
Clinical psych	yes (first used in '99)	yes

So stability looks like an important feature of a centralized labor market clearinghouse.

Formal Package Auctions Model

seller

players 0, 1, 2, 3, … J

bidders

- Bidders have:
  - quasi-linear utility
  - unrestricted budgets
  - full information
  - a finite set of packages of weakly positively valued packages of interest, including the empty set, which has zero value

Formal Package Auctions Model (cont.)

Seller has a feasible set,

$$X_0 \subseteq X_1 \times \cdots \times X_J \text{ with } (\emptyset, \ldots, \emptyset) \in X_0.$$

$$u_0 : X_0 \to \mathbb{R} \qquad u_0(\emptyset, \ldots, \emptyset) = 0$$

and a coalition, S, has a feasible assignment $\hat{x} \in F(S)$

if $\hat{x} \in X_0$ and for all j, if $j \notin S$ or $0 \notin S$, then $\hat{x}_j = \emptyset$

The pay off to a bidder, j, can be expressed as $\pi_j = u_j(x) - p_j$

Coalitional Value Function

$$w_u(S) = \max_{x \in F(S)} \sum_{j \in S} u_j(x_j)$$

Core of a cooperative game with transferable utility

$$Core(N, w) = \left\{ \pi \geqslant 0 \middle| \sum_{j \in N} \pi_j = w(N) \text{and} (\forall S \subseteq N) \sum_{j \in S} \pi_j \geqslant w(S) \right\}$$

Core and Strategyproofness

7

6

4

The outcome of a first price auction is in the core with respect to reported preferences, but the game is not strategyproof.

Strategyproof dictates how you should play. Core determines whether this is a game worth playing.

They’re grrrrrrreat!

PROPERTIES OF CORE SELECTION AUCTIONS

Sybilproof

Theorem 1 *An efficient direct auction mechanism has the property that no bidder can ever earn more than its Vickrey payoff by disaggregating and bidding with shills if and only if it is a core-selecting auction mechanism.*

	A	B	AB
X	8		
Z			10

→

	A	B	AB
X	8		
Y		8	
Z			10

Truncation

Theorem 2 *Suppose that* (f, P) *is a core-selecting direct auction mechanism and bidder* j *is favored. Let* $\hat{u}_{-j}$ *be any profile of reports of bidders other than* j*. Denote* j*'s actual value by* u_j *and let* $\bar{\pi}_j = w_{\hat{u}_{-j},u_j}(N) - w_{\hat{u}_{-j},u_j}(N-j)$ *be* j*'s corresponding Vickrey payoff. Then, the* $\bar{\pi}_j$ *truncation of* u_j *is among bidder* j*'s best replies in the mechanism and earns a payoff for* j *of* $\bar{\pi}_j$*. Moreover, this remains a best reply even in the expanded strategy space in which bidder* j *is free to use shills.*

7 → 6

6 6

4 4

Hold these fixed!

Bidder Optimal Mechanisms

Theorem 3 *For every valuation profile u and corresponding bidder optimal imputation π, the profile of π_j truncations of u_j is a full information equilibrium profile of every core selecting auction. The equilibrium goods assignment x^* maximizes the true total value $\sum_{i \in N} u_i(x_i)$, and the equilibrium payoff vector is π (including π_0 for the seller).*

	A	B	AB
X	8		
Y		2	
Z			10

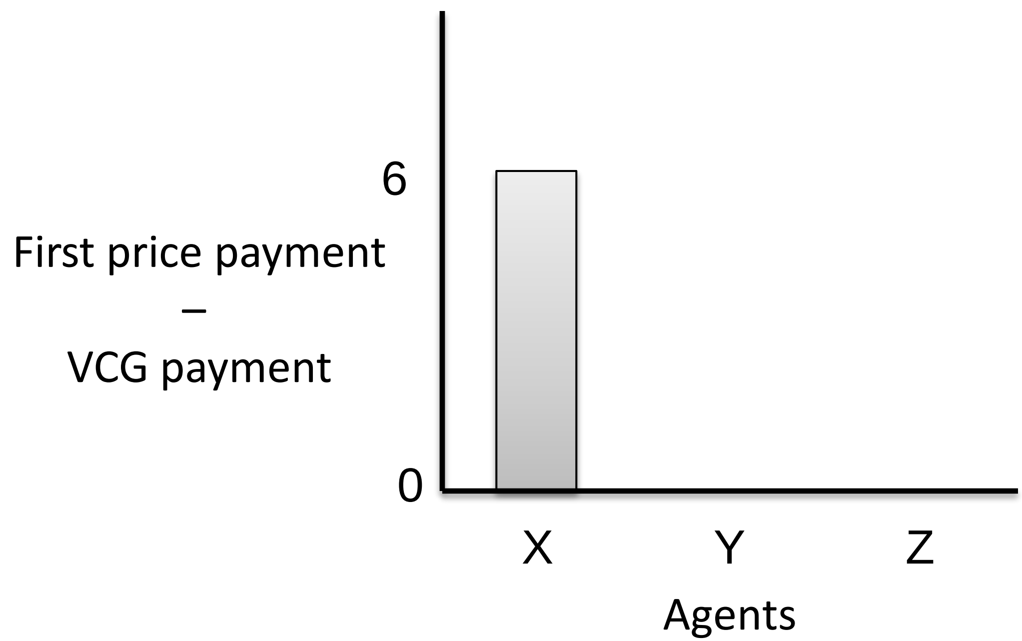

Bidder Optimal Mechanisms

Theorem 3 *For every valuation profile u and corresponding bidder optimal imputation π, the profile of π_j truncations of u_j is a full information equilibrium profile of every core selecting auction. The equilibrium goods assignment x^* maximizes the true total value $\sum_{i \in N} u_i(x_i)$, and the equilibrium payoff vector is π (including π_0 for the seller).*

	A	B	AB
X	7		
Y		3	
Z			10

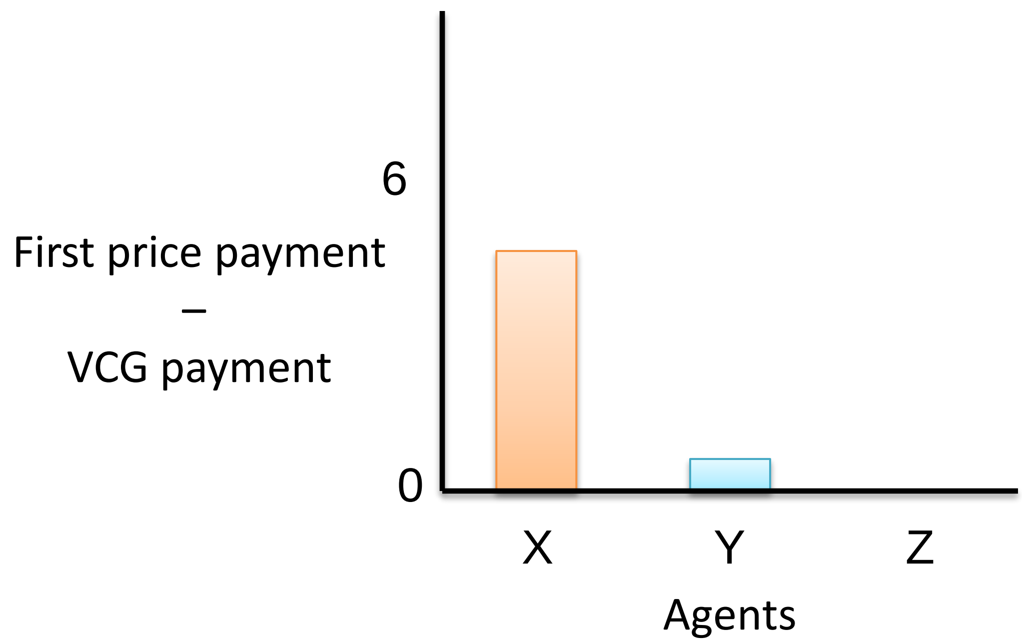

Bidder Optimal Mechanisms

Theorem 3 *For every valuation profile u and corresponding bidder optimal imputation π, the profile of π_j truncations of u_j is a full information equilibrium profile of every core selecting auction. The equilibrium goods assignment x^* maximizes the true total value $\sum_{i \in N} u_i(x_i)$, and the equilibrium payoff vector is π (including π_0 for the seller).*

	A	B	AB
X	6		
Y		4	
Z			10

Monotonicity of Revenues

Theorem 5 *The seller's minimum payoff in the core with bidder values $\hat{u}$ is non-decreasing in $\hat{u}$.*

- Core constraints are weakly stronger as bids increase
- Set of core allocations weakly shrinks as bids increase
- Thus, minimum payoff to seller over all core allocations weakly increases
  - Doesn't say anything about other core payoffs

Incentives and Regret

Definition The *incentive profile* for a core-selecting auction P at u is $\varepsilon^P = \left\{\varepsilon_j^P(u)\right\}_{j \in N-0}$ where $\varepsilon_j^P(u) \equiv \sup_{\hat{u}_j} u_j(f_j(u_{-j}, \hat{u}_j)) - P\left(u_{-j}, \hat{u}_j, f_j(u_{-j}, \hat{u}_j)\right)$ is j's maximum gain from deviating from truthful reporting when j is favored.

Theorem 4 *A core-selecting auction provides optimal incentives if and only if for every u it chooses a bidder optimal allocation.*

Corollary *When the Vickrey outcome is a core allocation, then truthful reporting is an ex post equilibrium for any mechanism that always selects bidder optimal core allocations.*

Properties of Core-Selecting Auctions

1. Sybilproof
2. In the full information game, each favored player has a bid which provides their VCG payoff
3. There exists a full-information Nash Equilibrium when the mechanism is bidder-optimal
4. Monotonicity of revenues for the seller
5. Incentives to misreport are minimized if and only if the mechanism is bidder-optimal

Connections to Marriage Problem

1. S-optimal stable matches are incentive compatible for S

Fred
- Wilma
- Betty

Barney
- Betty
- Wilma

Wilma
- Barney
- Fred

Betty
- Fred
- Barney

Connections to Marriage Problem

2. All deviations take the form of truncations

Fred
- Wilma
- Betty

Barney
- Betty
- Wilma

Wilma
- Barney
- ~~Fred~~

Betty
- Fred
- ~~Barney~~

Connections to Marriage Problem

3. Revenue monotonicity

Fred
- Wilma
- Betty

Barney
- Betty
- Wilma

Bam Bam
- Wilma
- Pebbles (not pictured)

Wilma
- Barney
- ~~Fred~~

Betty
- Fred
- ~~Barney~~

Critiques

- Equilibrium results only apply in full-information setting
- Theorem 2 relies on other bids being fixed
- Core constraints are enforced relative to reported preferences
- Incentives results are non-equilibrium

Lubin & Parkes

DECIDING A PAYMENT RULE

Payment rules should distribute all the surplus

Definition The *incentive profile* for a core-selecting auction P at u is $\varepsilon^P = \left\{\varepsilon_j^P(u)\right\}_{j \in N-0}$ where $\varepsilon_j^P(u) \equiv \sup_{\hat{u}_j} u_j(f_j(u_{-j}, \hat{u}_j)) - P\left(u_{-j}, \hat{u}_j, f_j(u_{-j}, \hat{u}_j)\right)$ is j's maximum gain from deviating from truthful reporting when j is favored.

Theorem 4 *A core-selecting auction provides optimal incentives if and only if for every u it chooses a bidder optimal allocation.*

Corollary *When the Vickrey outcome is a core allocation, then truthful reporting is an ex post equilibrium for any mechanism that always selects bidder optimal core allocations.*

Bidder Optimal Mechanisms

	A	B	AB
X	8		
Y		8	
Z			10

A Choice of Payments

Threshold

Allocate surplus to minimize the maximum $\Delta_{vcg,i} - \Delta_i$, subject to $\Delta_i \leq \Delta_{vcg,i}$, $\forall i \in \mathcal{N}$.

Small

Allocate surplus from smallest $\Delta_{vcg,i}$ to largest, never exceeding $\Delta_{vcg,i}$.

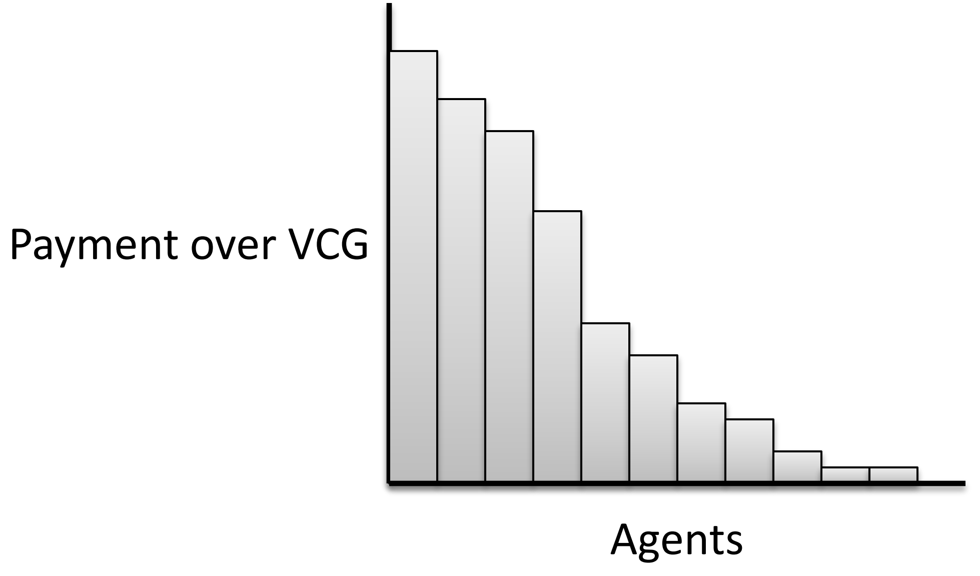

A Choice of Payments

Threshold

Allocate surplus to minimize the maximum $\Delta_{vcg,i}$ - Δ_i, subject to $\Delta_i \leq \Delta_{vcg,i}$, $\forall i \in \mathcal{N}$.

Small

Allocate surplus from smallest $\Delta_{vcg,i}$ to largest, never exceeding $\Delta_{vcg,i}$.

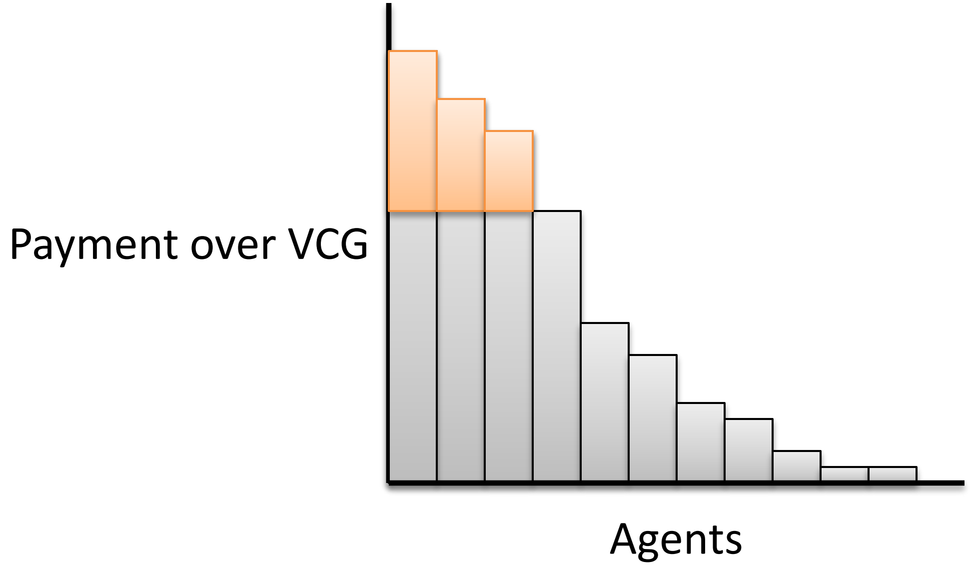

A Choice of Payments

Threshold

Allocate surplus to minimize the maximum $\Delta_{vcg,i} - \Delta_i$, subject to $\Delta_i \leq \Delta_{vcg,i}$, $\forall i \in \mathcal{N}$.

Small

Allocate surplus from smallest $\Delta_{vcg,i}$ to largest, never exceeding $\Delta_{vcg,i}$.

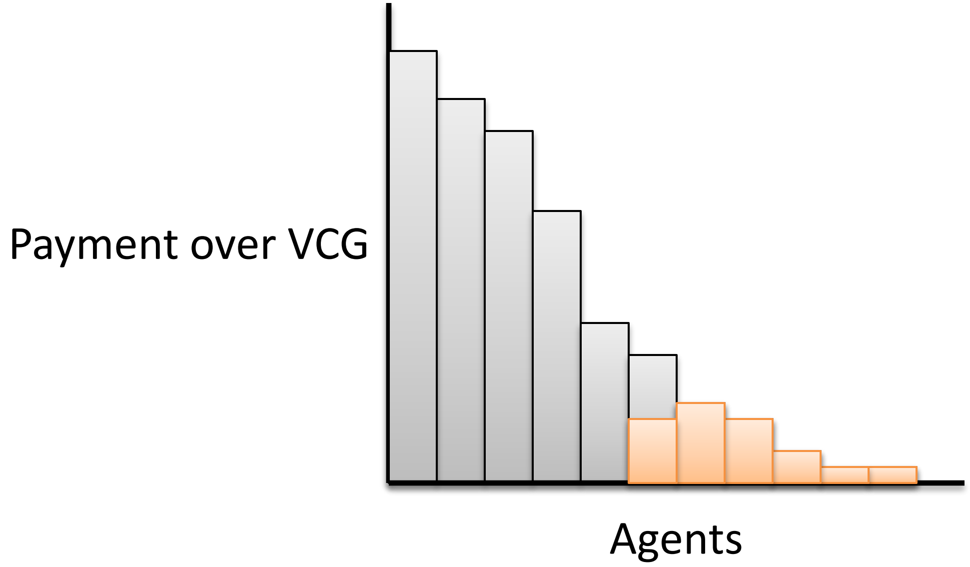

Distinctions from Milgrom’s package auctions

- Suitable for any setting, even games with no stable allocations
- Only applicable in settings with money
- Interested primarily in minimizing incentives to misreport given other constraints
  - Relative to a strategyproof reference (VCG)

AKA “Relative Entropy.” Not a distance or a metric but a state of mind

NORMALIZED KULLBACK-LIEBLER

Multivariate KL

$$\int_{\pi\in\Pi} H^*(\pi)\log\left(\frac{H^*(\pi)}{H^m(\pi)}\right)d\pi.$$

Where $m = (f, p)$ is the mechanism under consideration, comprised of an allocation and payment rule. m^* is our strategyproof reference mechanism (VCG payments).

$\pi^m(v) = (\pi_1(v), \ldots, \pi_n(v))$ and $\pi^*(v) = (\pi_1^*(v), \ldots, \pi_n^*(v))$ are the payoff vectors for an instance of the mechanism and reference mechanism, respectively.

$\pi_i(v) = v_i(f(v)) - p_i(v)$ is the payoff to an agent, i.

$H^m(\pi)$ and $H^*(\pi)$ are distributions over the payoff vectors of the mechanism and reference mechanism, respectively, as induced by some distribution on valuations.

One dimensional KL

payoff vectors under m

(0, 9, 1, 3, 0, 0, 0)
(2, 7, 3, 0, 0, 0, 1)
(4, 4, 0, 0, 0, 5, 7)

→ (0, 9, 1, 3, 0, 0, 0, 2, 7, 3, 0, 0, 0, 1, 4, 4, 0, 0, 0, 5, 7)

payoff vectors under m*

(0, 4, 1, 2, 0, 0, 0)
(2, 4, 1, 0, 0, 0, 1)
(4, 2, 0, 0, 0, 5, 4)

Restriction to active agents

payoff vectors under m

(0, 9, 1, 3, 0, 0, 0)
(2, 7, 3, 0, 0, 0, 1)
(4, 4, 0, 0, 0, 5, 7)

→ (0, 9, 1, 3, 0, 0, 0, 2, 7, 3, 0, 0, 0, 1, 4, 4, 0, 0, 0, 5, 7)

↓

(9, 1, 3, 2, 7, 3, 1, 4, 4, 5, 7)

payoff vectors under m*

(0, 4, 1, 2, 0, 0, 0)
(2, 4, 1, 0, 0, 0, 1)
(4, 2, 0, 0, 0, 5, 4)

Normalization

payoff vectors under m

(0, 9, 1, 3, 0, 0, 0)
(2, 7, 3, 0, 0, 0, 1)
(4, 4, 0, 0, 0, 5, 7)

→ (0, 9, 1, 3, 0, 0, 0, 2, 7, 3, 0, 0, 0, 1, 4, 4, 0, 0, 0, 5, 7)

↓

(9, 1, 3, 2, 7, 3, 1, 4, 4, 5, 7)

↓

payoff vectors under m*

(0, 4, 1, 2, 0, 0, 0)
(2, 4, 1, 0, 0, 0, 1)
(4, 2, 0, 0, 0, 5, 4)

→ (9/6, 1/6, 3/6, 2/5, 7/5, 3/5, 1/5, 4/4, 4/4, 5/4, 7/4)

Constructing a distribution

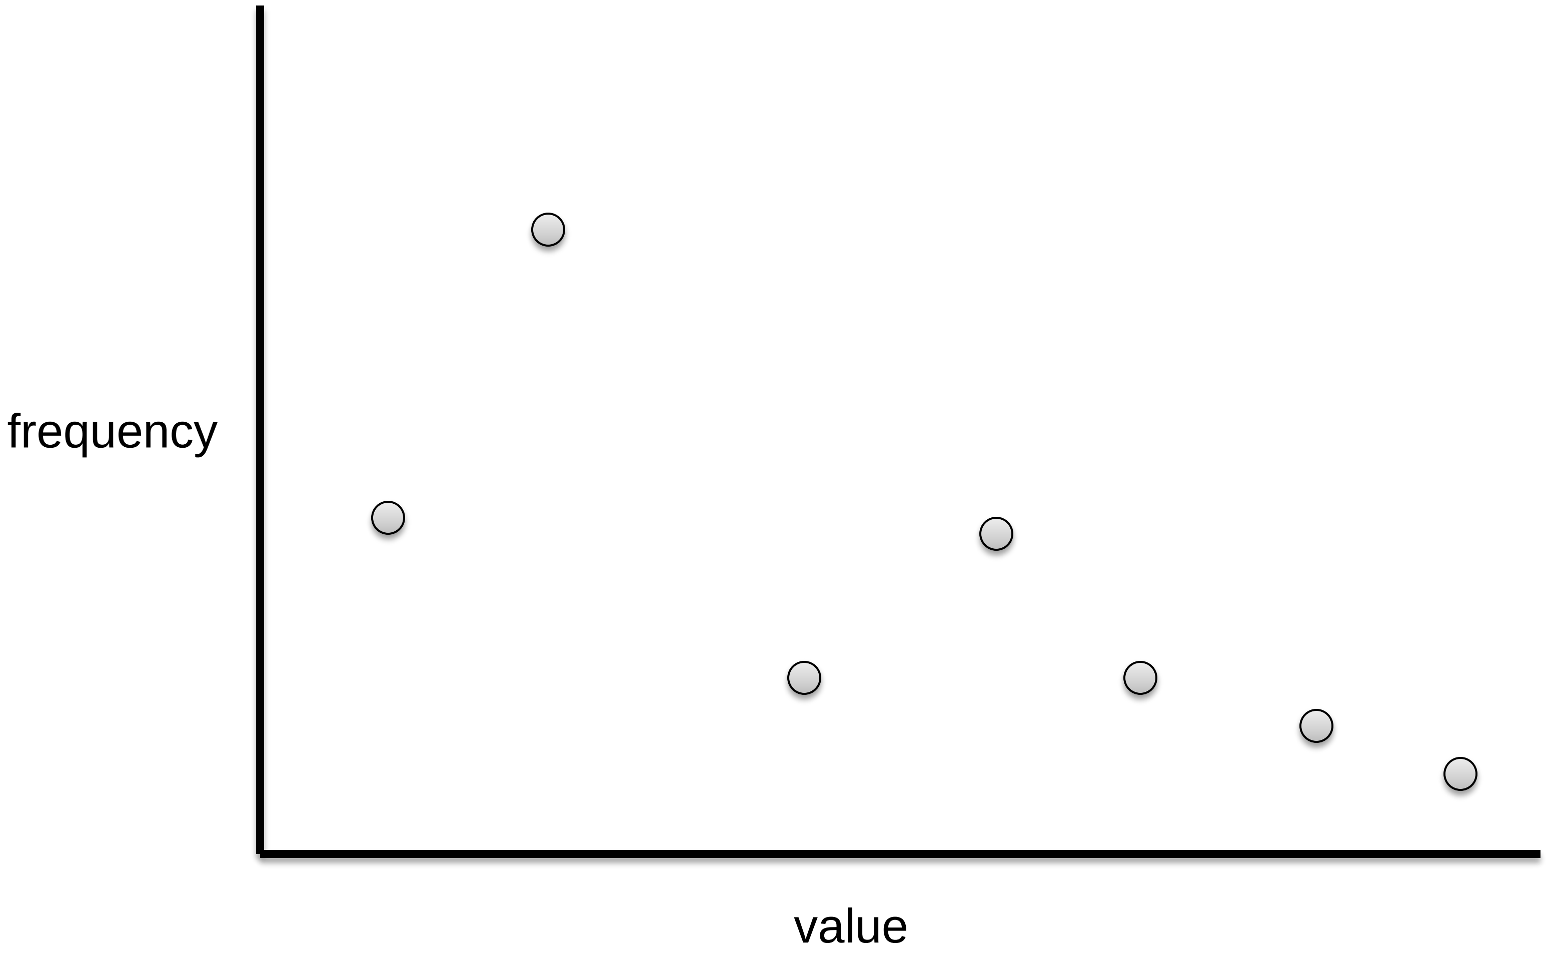

Constructing a distribution

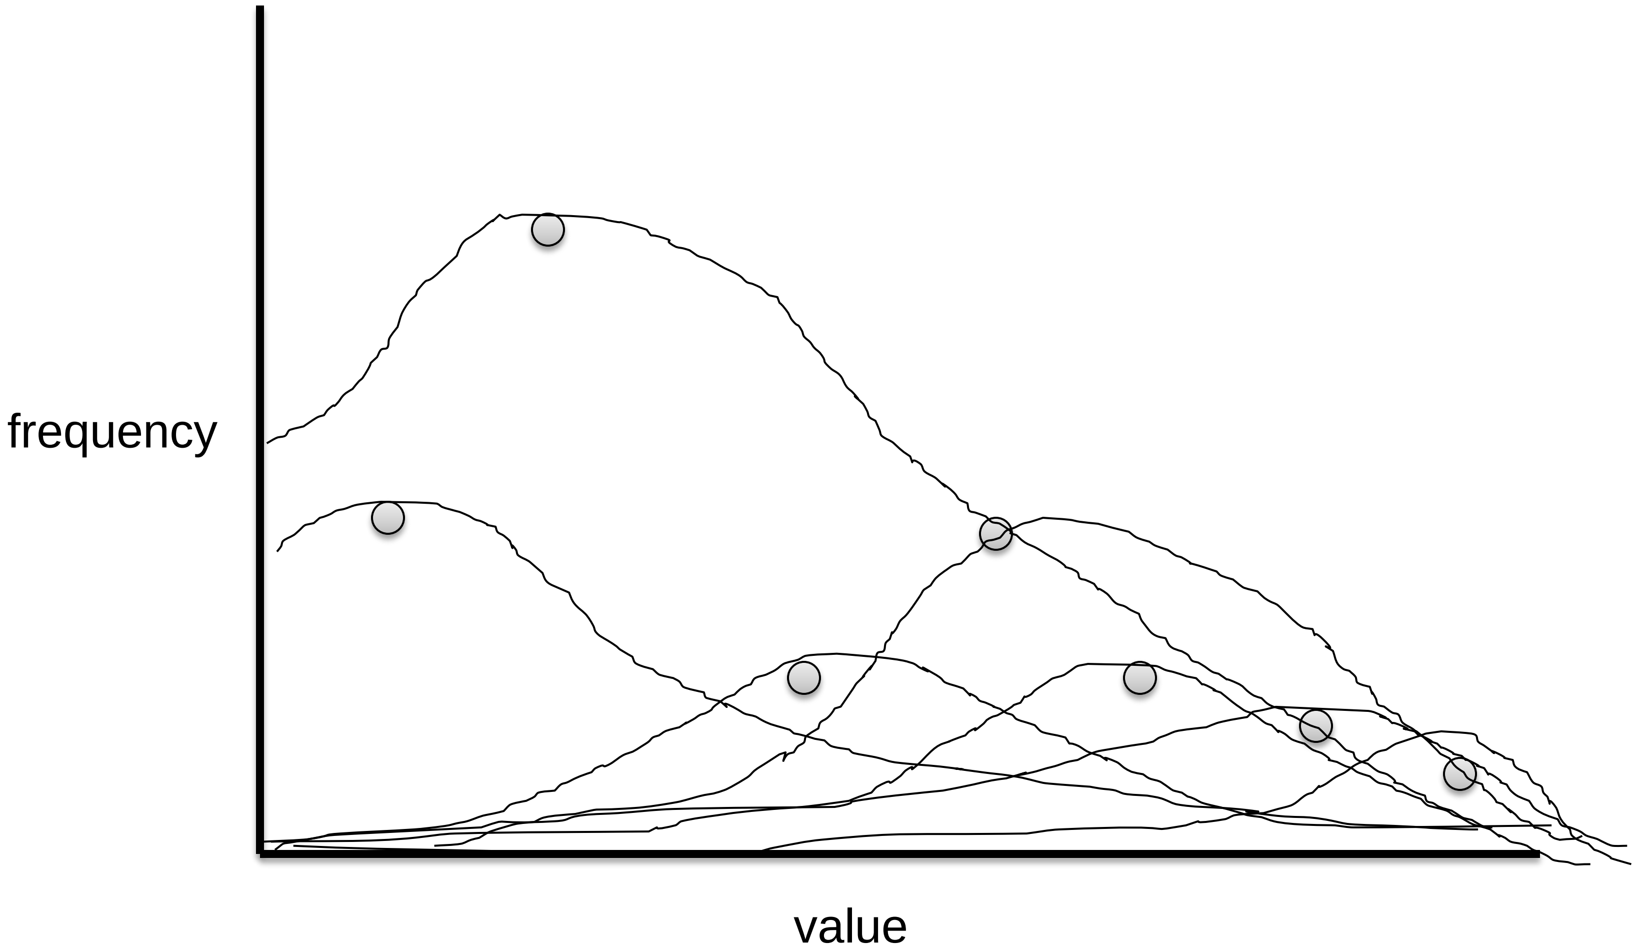

Constructing a distribution

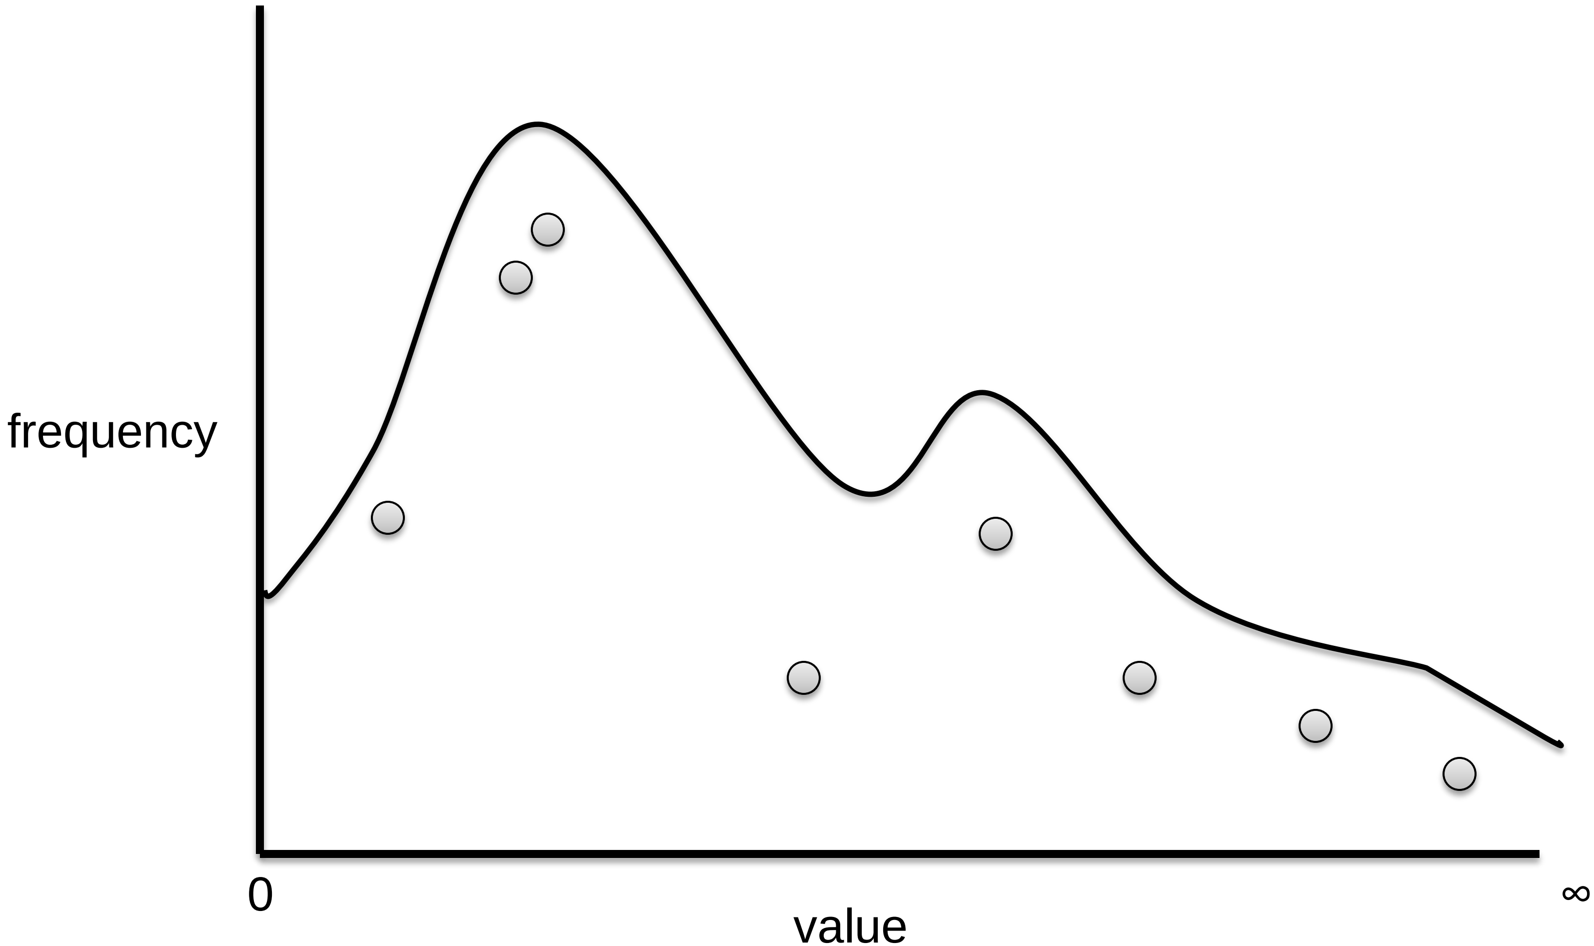

Normalized Kullback-Leibler

$$KLnorm(m) = \int_0^\infty \widehat{H}^*(\pi) log \left(\frac{\widehat{H}^*(\pi)}{\widehat{H}^m(\pi)} \right) d\pi$$

KL in action

EMPIRICALLY EVALUATING PAYMENT RULES

Testing a metric

1. Consider a set of payment rules
   – Small, Threshold, etc.
2. Model the agents
   – Decay, Uniform and Super
3. Model equilibrium
   – Focus on a particular class of equilibrium that can be computed

Empirical Setup

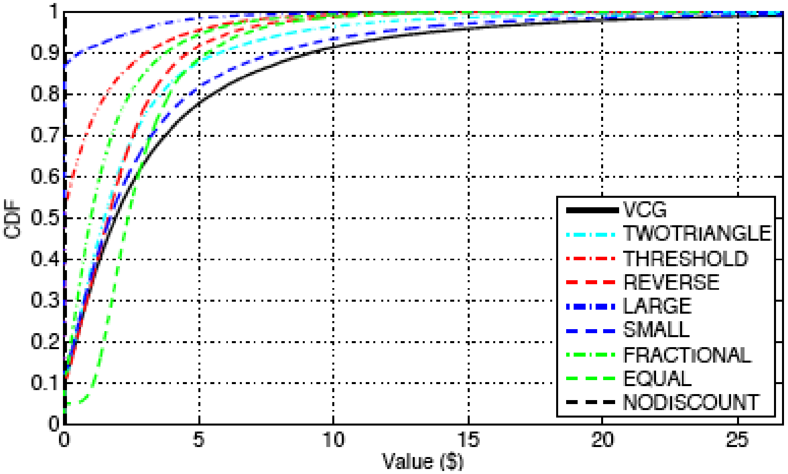

Figure 3: Distribution of payoffs in each mechanism

Mechanism	$KL norm$	$L_1\ norm$	$L_2\ norm$	$L_\infty\ norm$
Two Triangle	0.0735	**0.5914**	0.3170	0.1917
Threshold	0.0472	**0.5914**	**0.2355**	**0.1016**
Reverse	0.1251	**0.5914**	0.3066	0.2210
Small	**0.0452**	**0.5914**	0.4208	0.3527
Large	0.0559	**0.5914**	0.3110	0.2070
Fractional	0.0741	**0.5914**	0.2528	0.1513
Equal	0.3043	0.8037	0.3727	0.2576
No Discount	0.6372	1.5876	0.6679	0.4030

Table 1: Metric value at truth averaged across all three CE scenarios. Minimal metric values in **bold.**

Equilibrium Computation

- Intractable to compute a Bayes-Nash Equilibrium for every instance of the CE
- Restrict attention to a specific class of Equilibrium that can be found numerically
  - Every player uses some fixed shave factor a
  - Bidders report (1-a)v
  - Sellers report (1+a)v
  - Can also use multiple factors a_1, a_2, a_3
  - Iterative method for optimizing shave factors to find candidate equilibrium

Evaluating Mechanisms

	One Equilibrium Class						Three Equilibrium Classes					
	Shave Factor			**Efficiency (%)**			**Shave Factor**			**Efficiency (%)**		
Rule	**Dec.**	**Uni.**	**Sup.**	**Dec.**	**Uni.**	**Sup.**	**Dec.**	**Uni.**	**Sup.**	**Dec.**	**Uni.**	**Sup.**
VCG	*0.0*	*0.0*	*0.0*	*100*	*100*	*100*	*0.0*	*0.0*	*0.0*	*100*	*100*	*100*
Two Triangle	**0.1**	**0.2**	**0.6**	**99.99**	**100**	**99.99**	**0.1**	**0.4**	**5.6**	**99.99**	**100**	**97.95**
Threshold	12.0	28.7	10.7	99.09	97.43	98.01	14.6	27.2	11.2	93.64	81.09	89.74
Reverse	14.9	57.7	52.3	98.70	83.38	51.52	13.0	65.8	57.6	98.99	77.30	56.08
Small	**0.1**	**0.2**	**0.3**	**99.99**	**100**	**100**	**0.0**	**0.1**	**0.2**	**99.99**	**100**	**100**
Large	**2.6**	**2.3**	9.8	**99.96**	**99.99**	98.26	**2.8**	**2.9**	67.1	**99.96**	**99.98**	78.83
Fractional	71.2	71.1	53.0	59.39	67.34	49.07	62.7	81.9	62.0	37.12	63.09	56.77
Equal	75.4	77.6	52.5	51.96	55.76	51.01	62.2	78.3	66.8	33.35	54.21	52.19
No Discount	75.6	76.0	53.2	51.56	59.01	48.23	62.3	80.9	72.4	34.15	50.11	48.21

Table 2: Restricted Bayes-Nash equilibrium: Shave Factor and Allocative Efficiency in Each Mechanism.

Evaluating Mechanisms

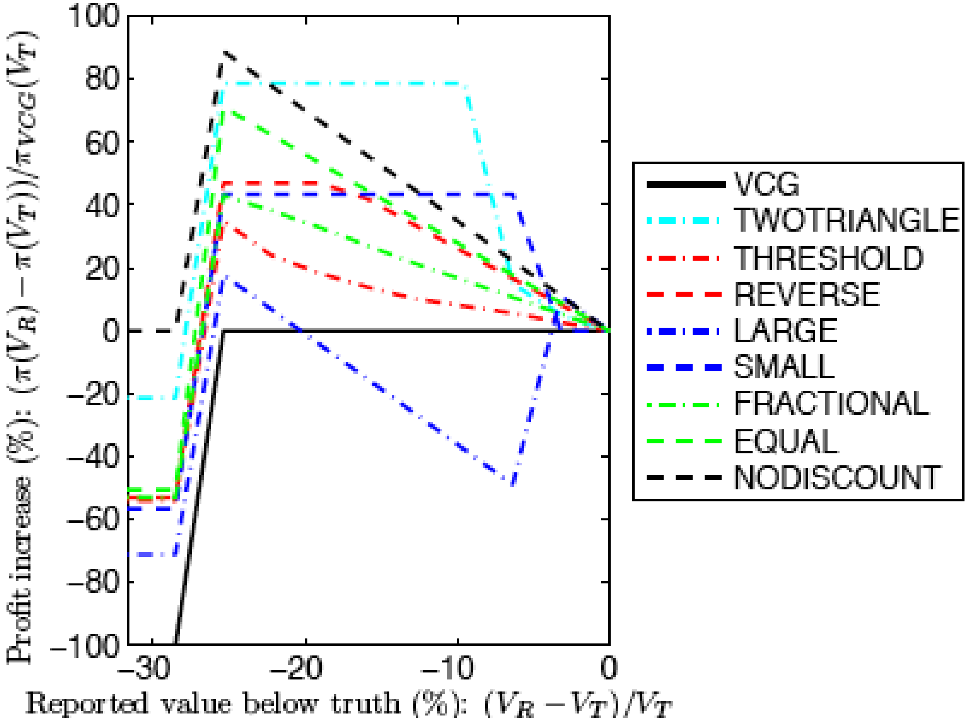

Figure 4: Profit gain by unilateral mis-report.

Evaluating Mechanisms: Threshold

Threshold

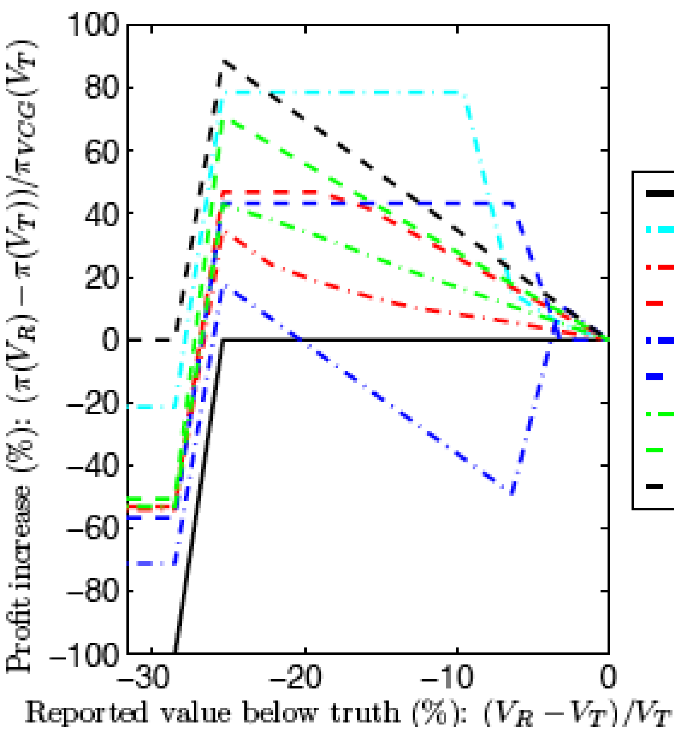

Figure 4: Profit gain by unilateral mis-report.

Allocate surplus to minimize the maximum $\Delta_{vcg,i} - \Delta_i$, subject to $\Delta_i \leq \Delta_{vcg,i}$, $i \in \mathcal{N}$.

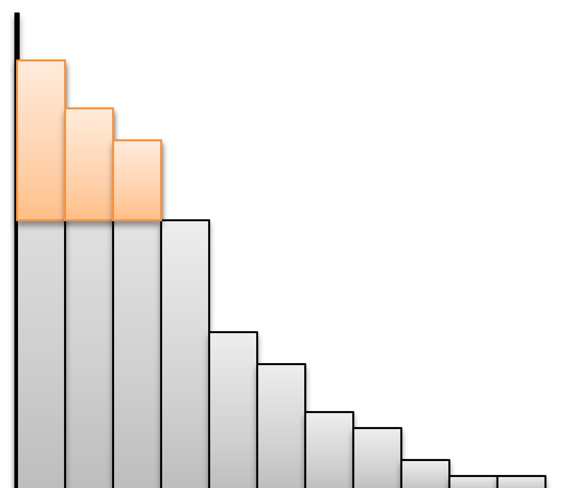

Evaluating Mechanisms: Small

Small

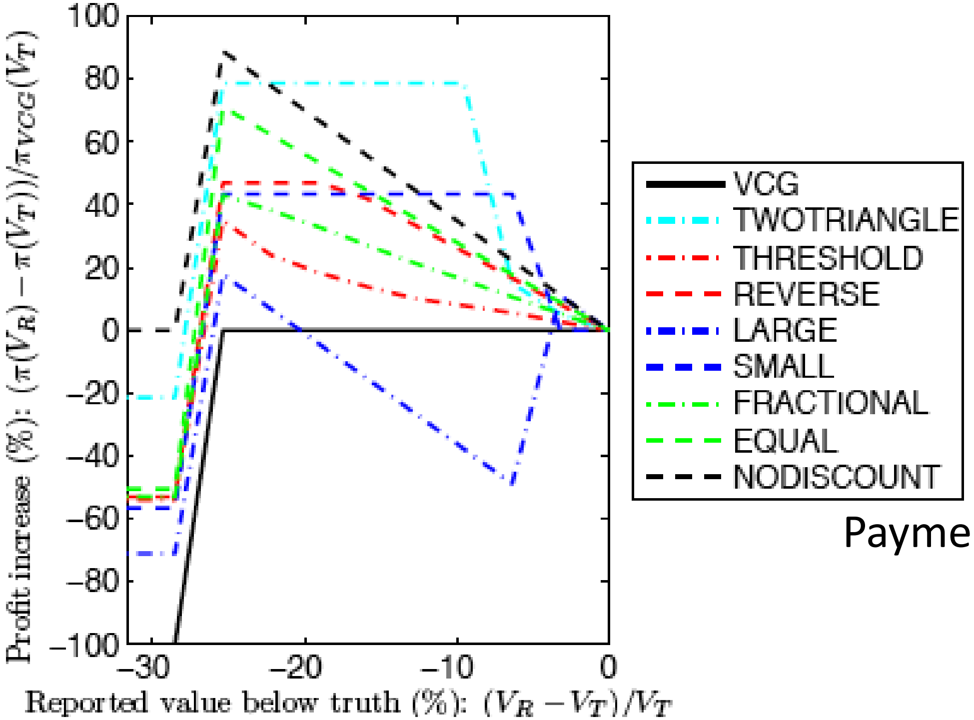

Figure 4: Profit gain by unilateral mis-report.

Allocate surplus from smallest $\Delta_{vcg,i}$ to largest, never exceeding $\Delta_{vcg,i}$.

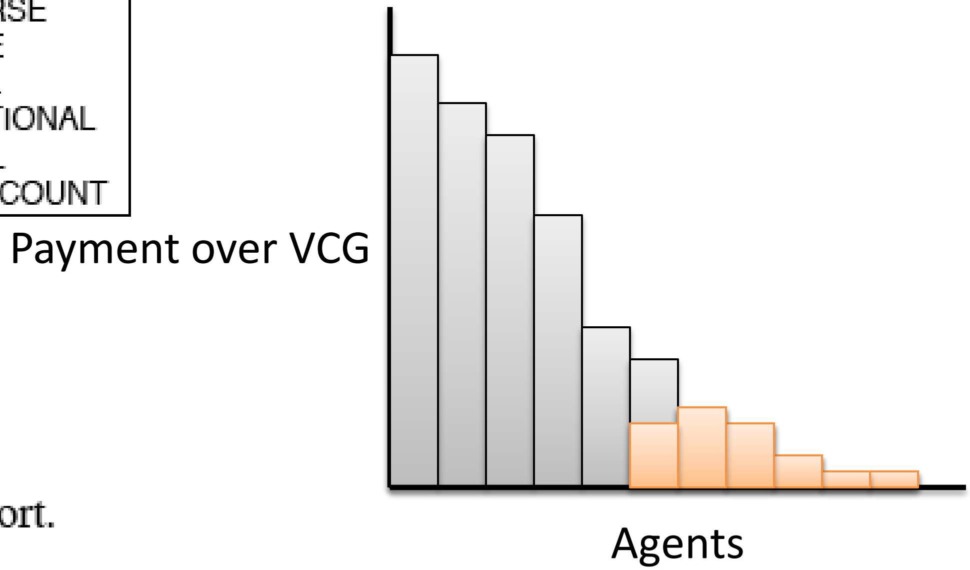

Evaluating Metrics

Correlation with Efficiency at Truth			
Metric	**Corr.**	**ρ-value**	**Significant?**
$KLnorm$	-0.3814	0.0044	Y
$L_1\,norm$	-0.1698	0.2197	N
$L_2\,norm$	0.0154	0.9120	N
$L_\infty\,norm$	0.0220	0.8745	N

Correlation with Mean Shave at Truth			
Metric	**Corr.**	**ρ-value**	**Significant?**
$KLnorm$	0.3794	0.0047	Y
$L_1\,norm$	0.1610	0.2447	N
$L_2\,norm$	-0.1001	0.4712	N
$L_\infty\,norm$	-0.1147	0.4087	N

Table 3: Correlation between metrics evaluated at truth and both efficiency and the amount of shaving, considering all 54 conditions (Significance at 0.05 level)

Correlation with Efficiency in Equilibrium			
Metric	**Corr.**	**ρ-value**	**Significant?**
$KLnorm$	-0.4989	1.2292e-04	Y
$L_1\,norm$	-0.6460	1.3269e-07	Y
$L_2\,norm$	-0.5119	7.6150e-05	Y
$L_\infty\,norm$	-0.3762	0.0051	Y

Correlation with Mean Shave in Equilibrium			
Metric	**Corr.**	**ρ-value**	**Significant?**
$KLnorm$	0.2702	0.0482	Y
$L_1\,norm$	0.5870	3.0820e-06	Y
$L_2\,norm$	0.4615	4.4464e-04	Y
$L_\infty\,norm$	0.3738	0.0054	Y

Table 4: Correlation between metrics evaluated at equilibrium and both the efficiency and the amount of shaving, considering all 54 conditions (Significance at 0.05 level)

Evaluating Metrics

Mechanism	$KL norm$	$L_1 norm$	$L_2 norm$	$L_\infty norm$
Two Triangle	0.0820	0.6096	0.3271	0.1976
Threshold	0.0556	0.6991	**0.2984**	**0.1367**
Reverse	0.1421	0.9415	0.4896	0.3104
Small	**0.0452**	**0.5903**	0.4208	0.3534
Large	0.0668	0.8269	0.4494	0.2916
Fractional	0.1303	1.1456	0.5683	0.3477
Equal	0.2033	1.3758	0.7291	0.4919
No Discount	0.3114	1.9962	1.0311	0.6721

Table 5: Metric value at equil. averaged across all three scenarios and equil. classes. Minimal values in **bold.**

Online Mechanism Selection

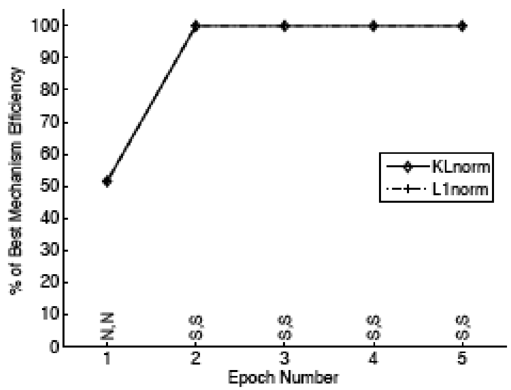

(a) Decay, 1 Class

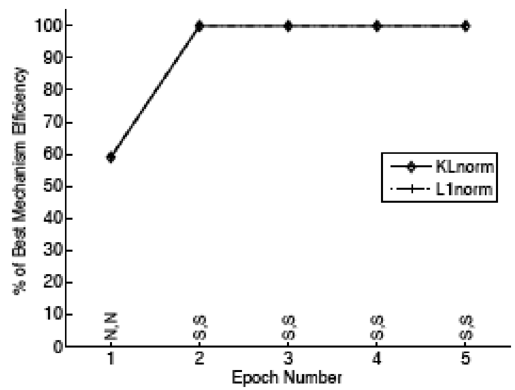

(b) Uniform, 1 Class

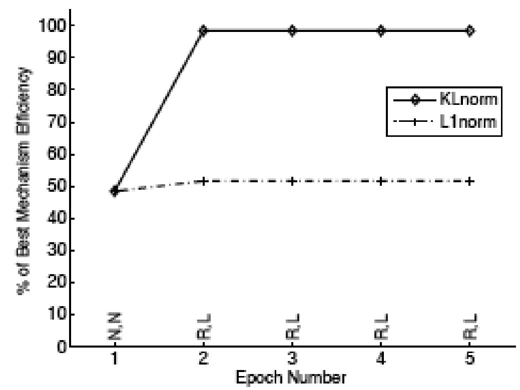

(c) Super, 1 Class

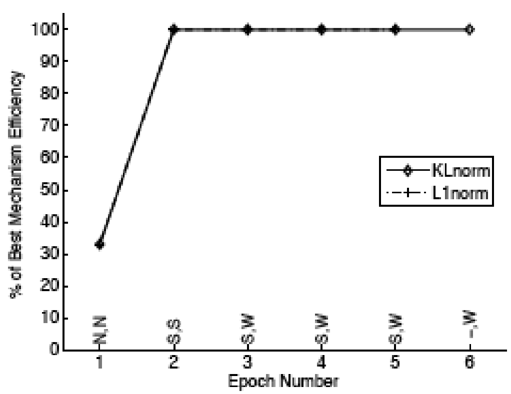

(d) Decay, 3 Class

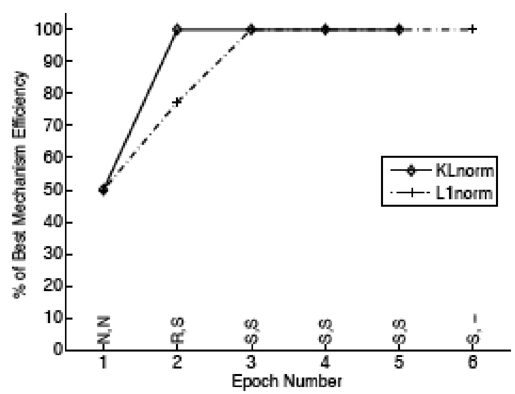

(e) Uniform, 3 Class

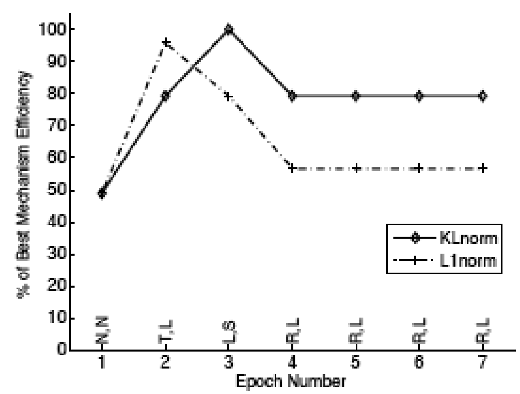

(f) Super, 3 Class

"By how much do the incentives for truthful reporting fail when other design objectives are imposed as constraints?"

IN CONCLUSION

Conclusion: Lubin & Parkes

- Two mechanisms are similar if their distribution over outcomes is similar
  - Outcomes are the observable of the mechanism, how you learn to play
- But the KL metric may not be the best
  - This was not an optimization question; it was justification of the KL metric and inspection of mechanisms' payoff distributions

Conclusion: Day & Milgrom

- Core stability is often consider a necessary condition in matching, even though SP is not
- Bidder-optimal core payments allow many necessary properties that VCG doesn’t have, while minimizing incentives to misreport