Combinatorial Exchanges

David C. Parkes
Harvard University

- What is a combinatorial exchange?

- Two-sided
- Complex valuations (swaps, contingent swaps, all-or-nothing sells, etc.)

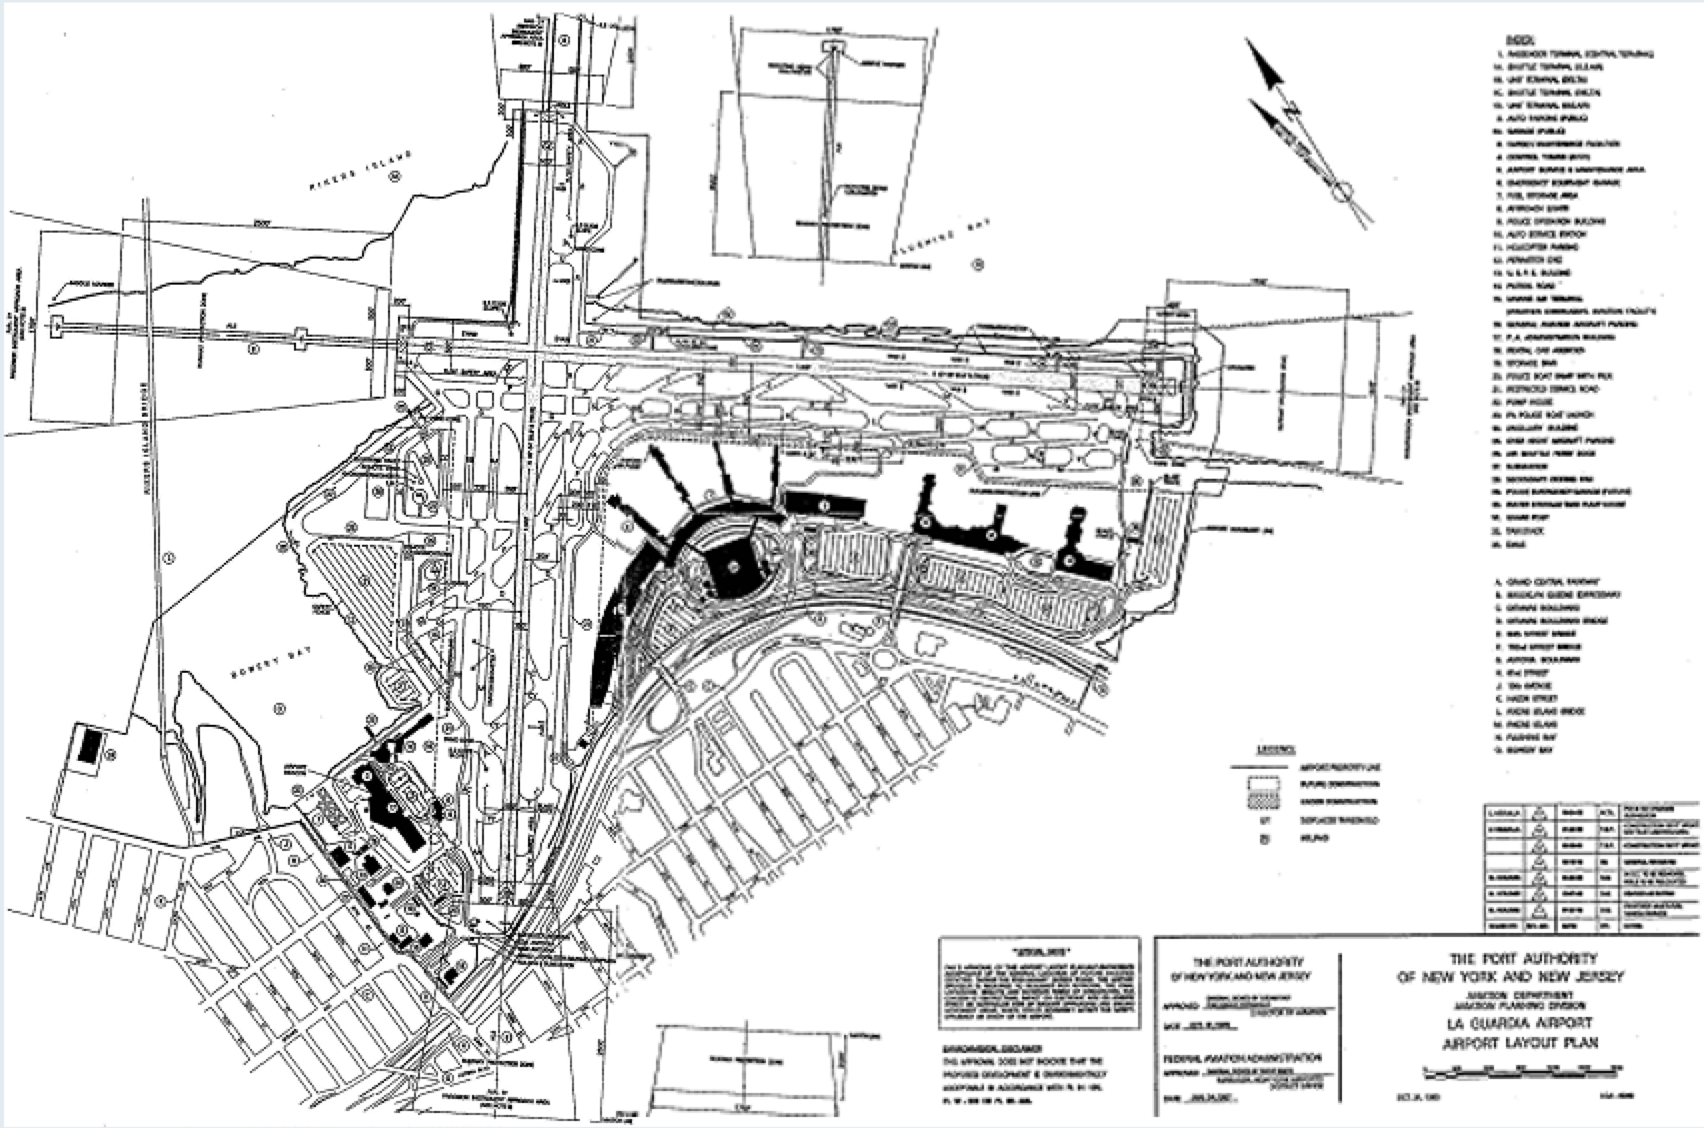
THE PORT AUTHORITY
OF NEW YORK AND NEW JERSEY
FEDERAL AVIATION ADMINISTRATION
THE PORT AUTHORITY
OF NEW YORK AND NEW JERSEY
LA GUARDIA AIRPORT
AIRPORT LAYOUT PLAN

Fragmented Spectrum

(E.Kwerel)

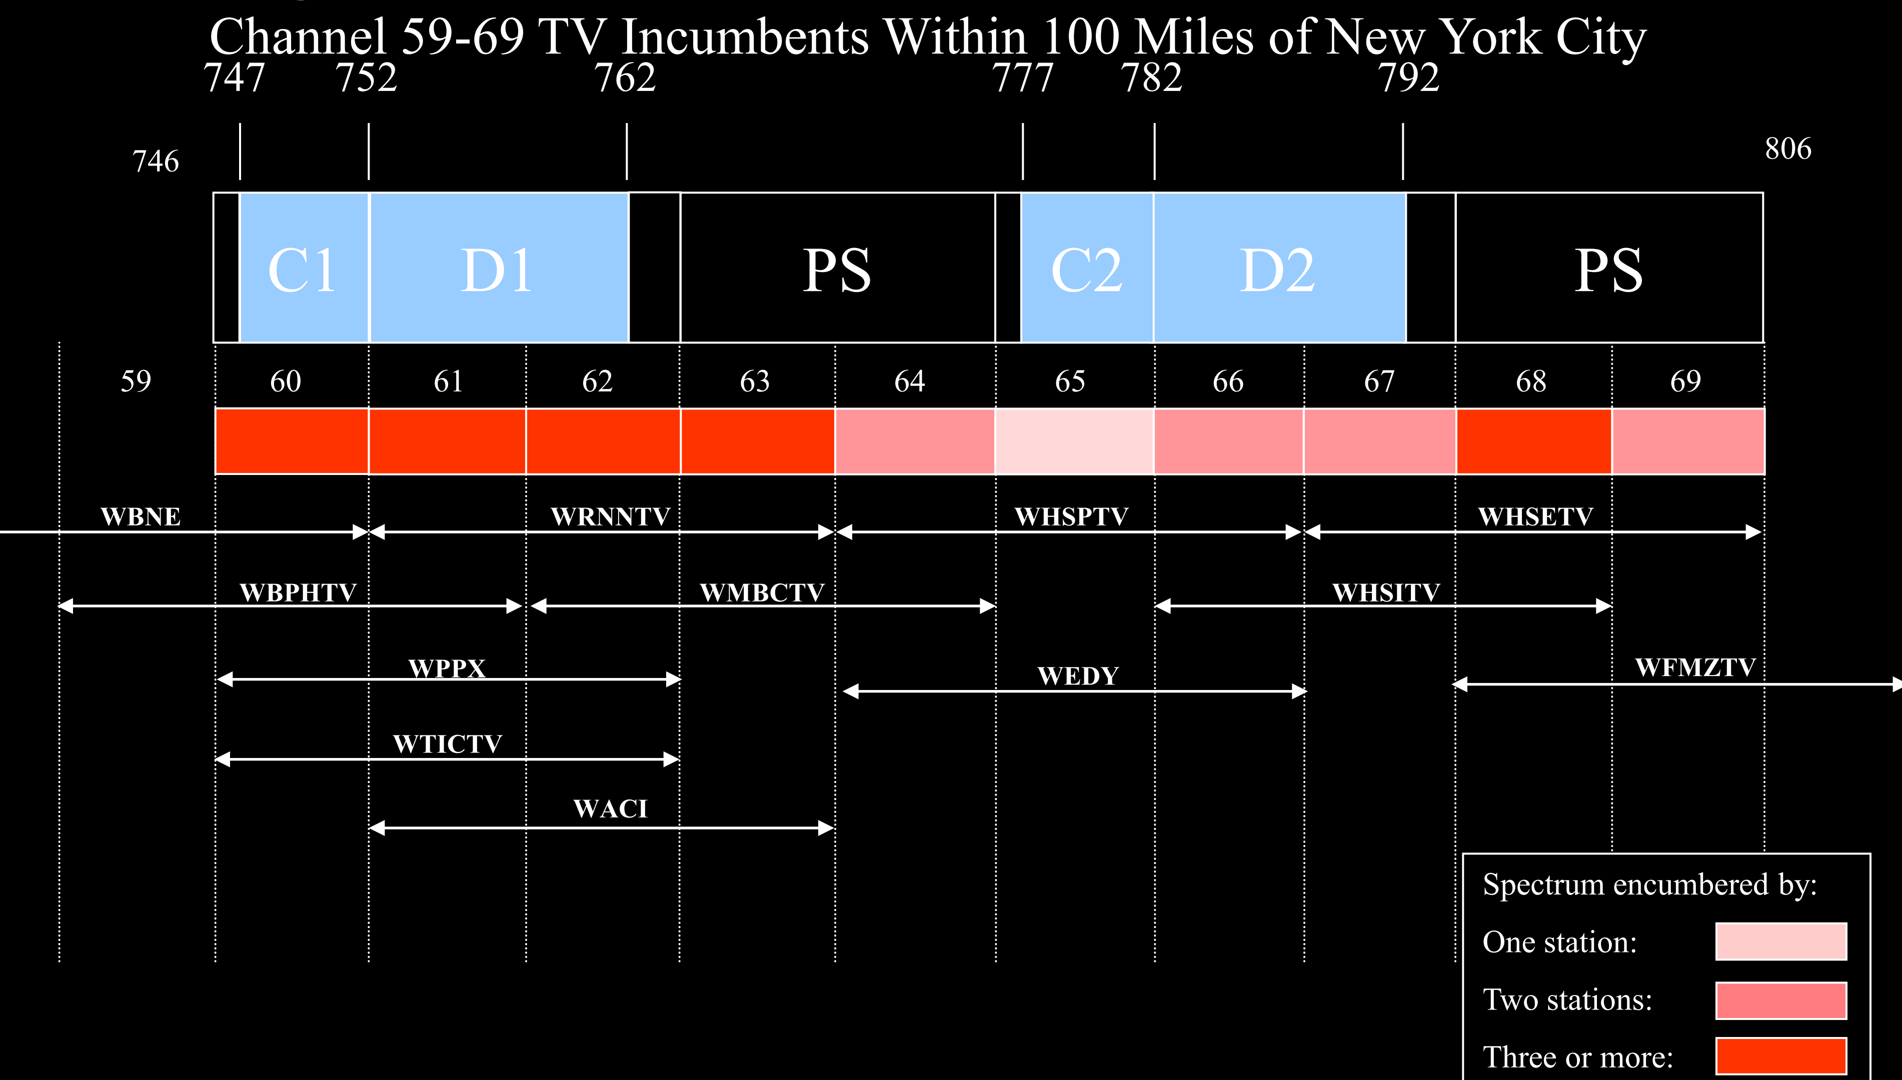

Challenges

- Expressive bidding language
- Preference Elicitation / Price discovery
- Scalable Winner Determination
- Payments:
  - Incentive compatibility
  - Stability
  - Fairness
  - …

Computational MD

- Economic constraints
  - e.g., incentive compatibility, core, etc.
- Computational constraints
  - e.g., scalable winner determination, minimal preference elicitation, etc.

“Mechanism” = “Algorithm”

Exchange Example

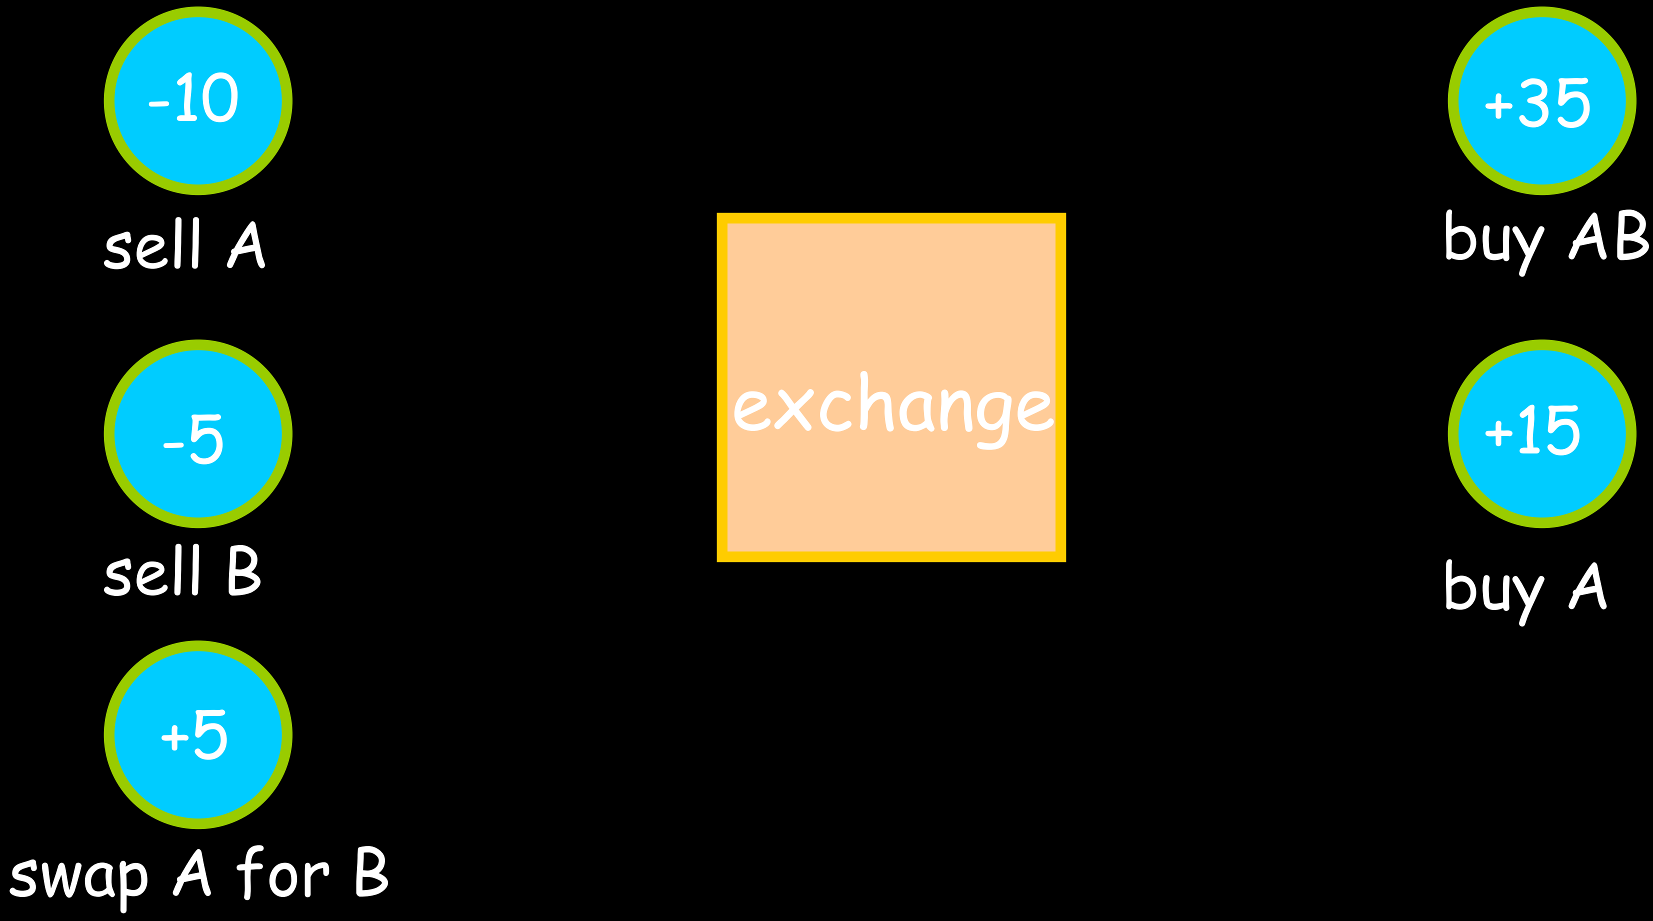

Exchange Example

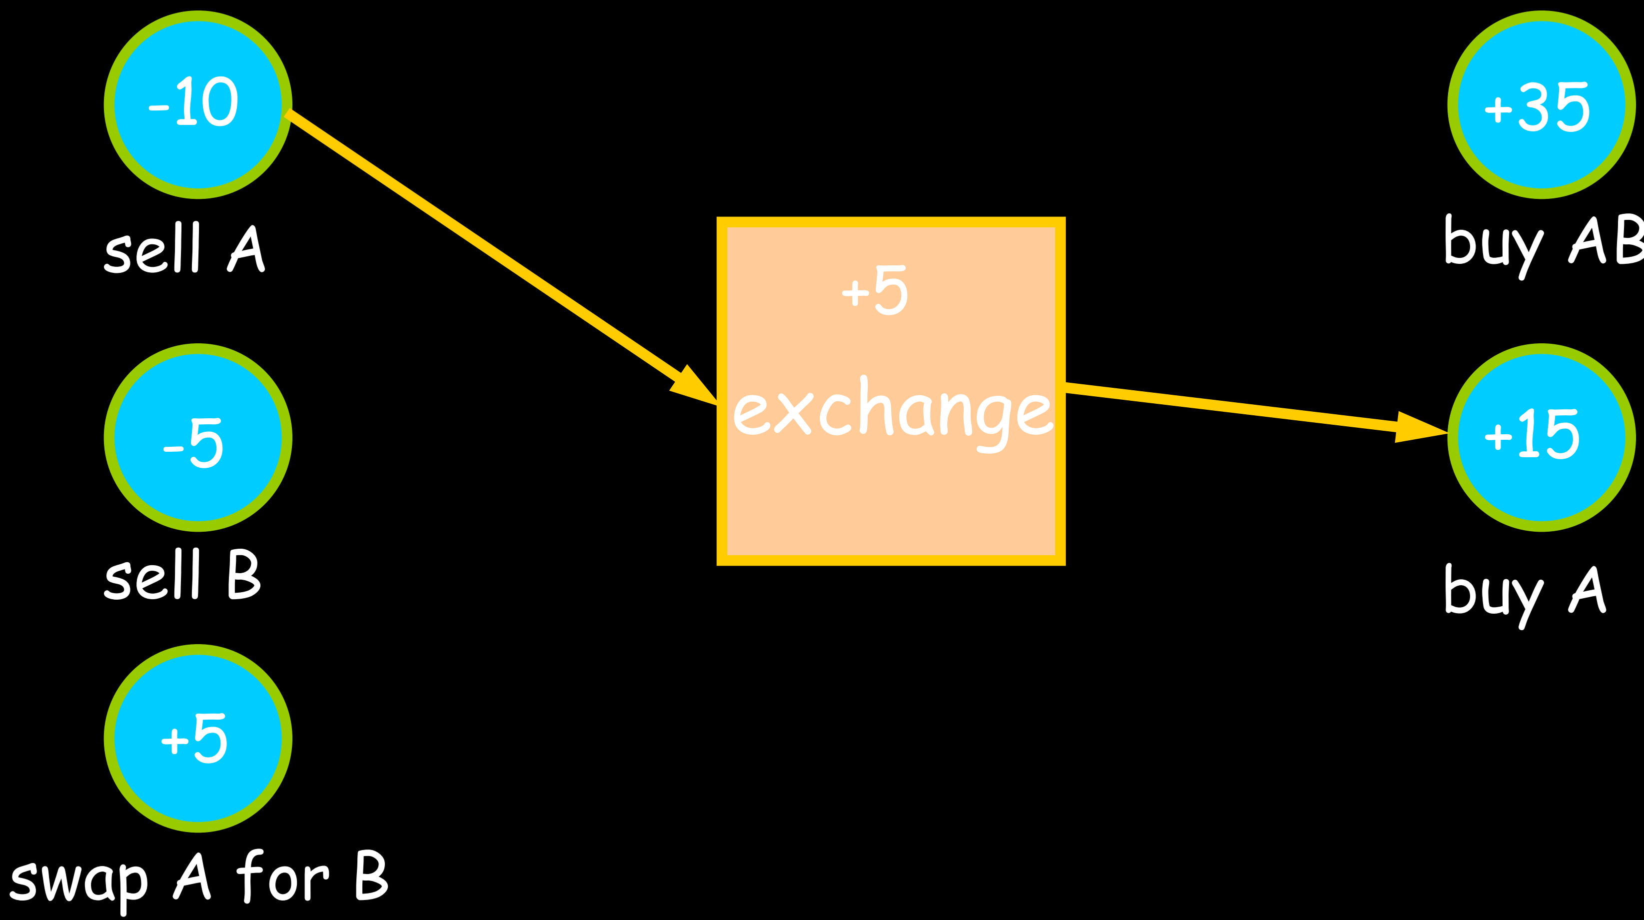

Exchange Example

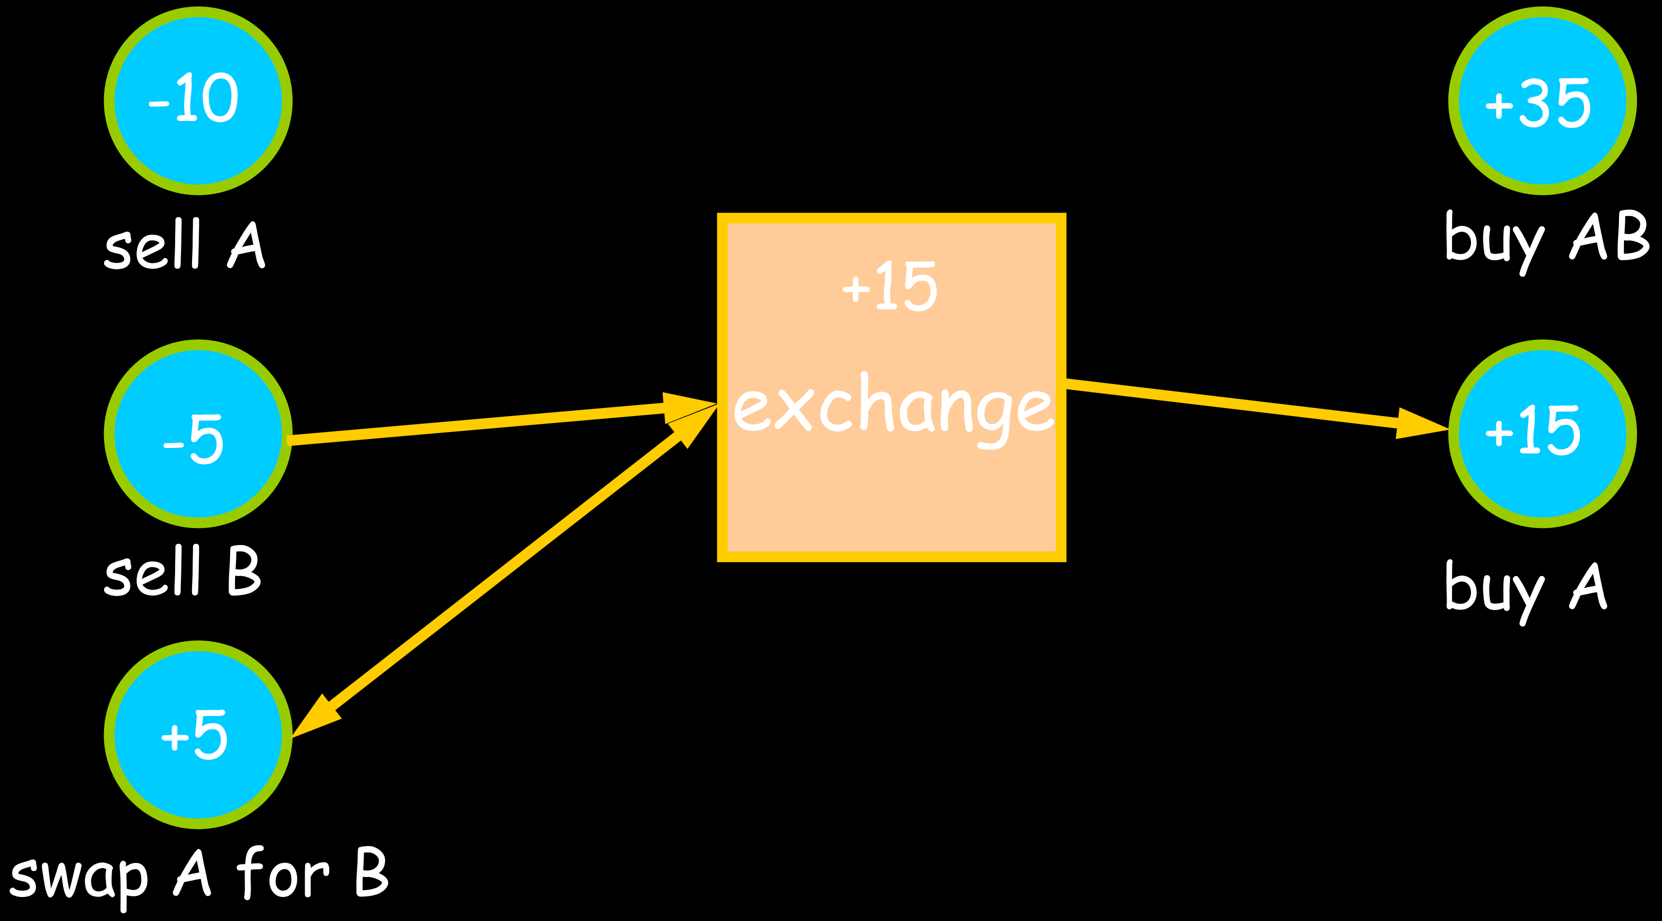

Exchange Example

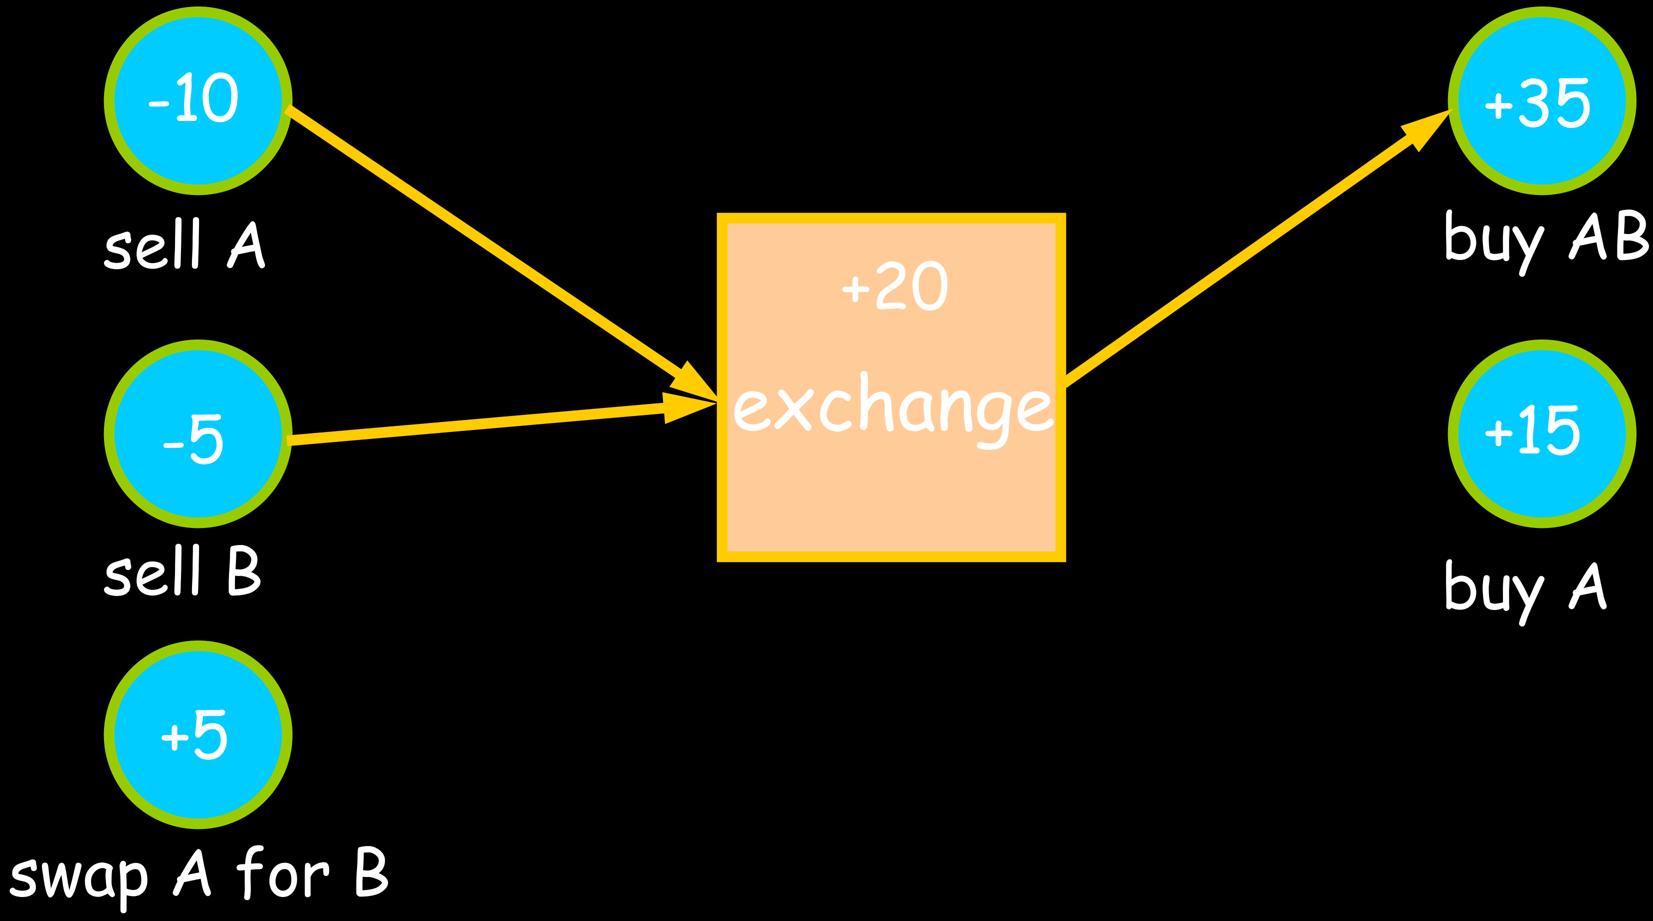

Exchange Example

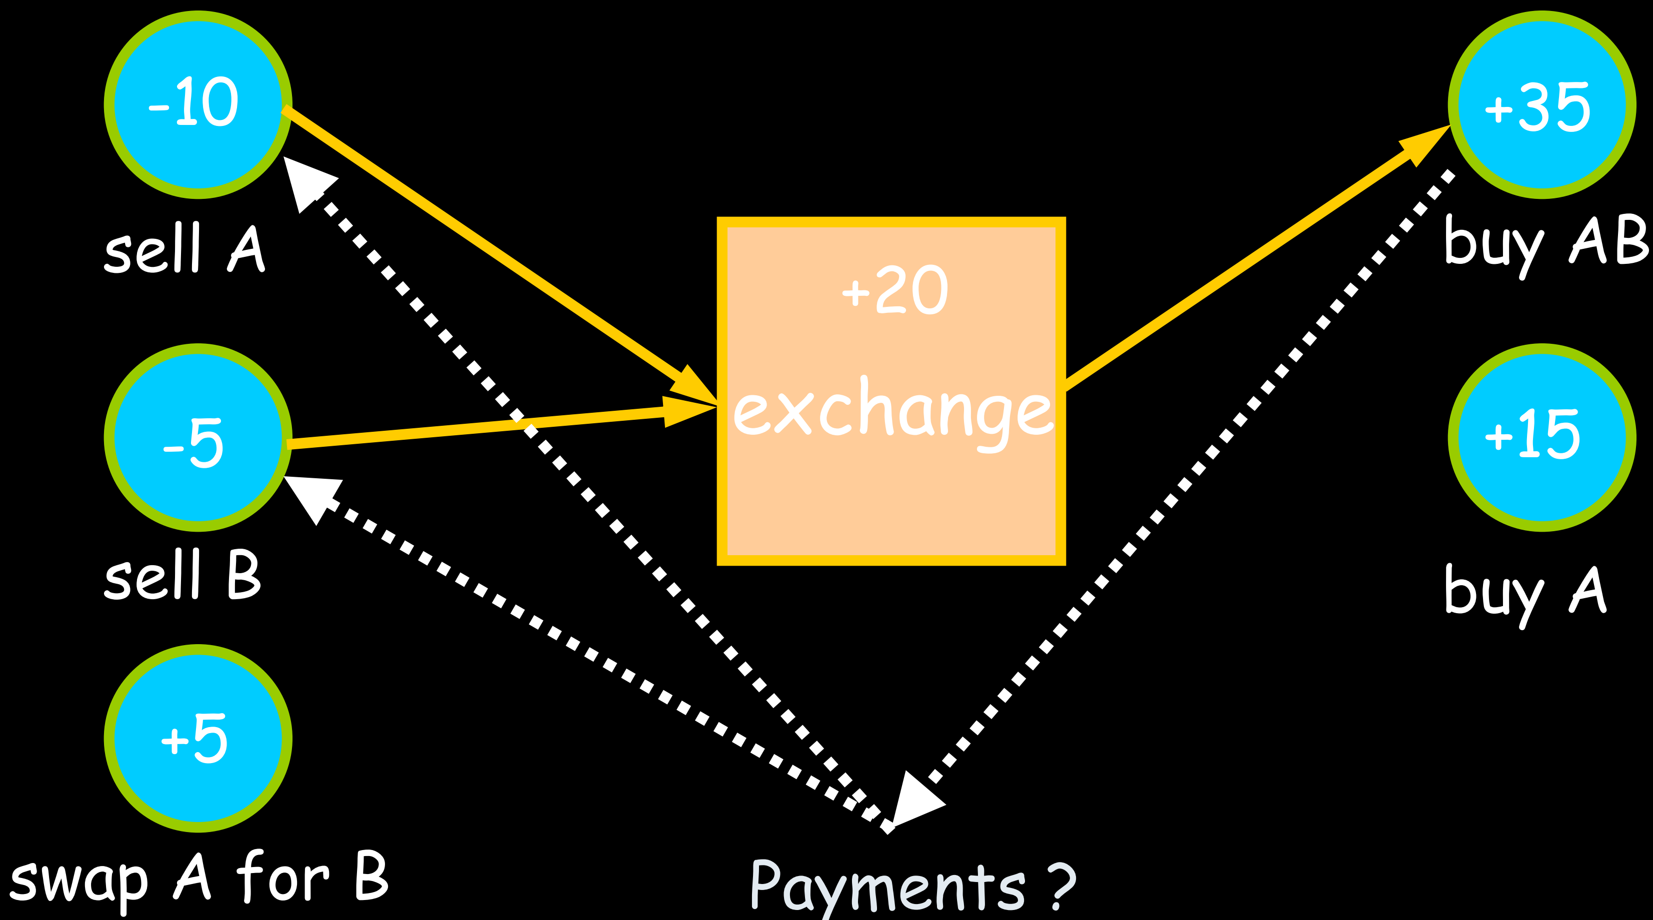

Tree Based Bidding Language

(Cavallo et al.'05)

- Defines change in value for a trade
  - entirely symmetric for buyers and sellers
  - "sell AB, value -\$100"; "buy A, value +\$20"

- Generalizes XOR, OR, XOR/OR

(Sandholm'99, Nisan'00)

Example 1: “and”

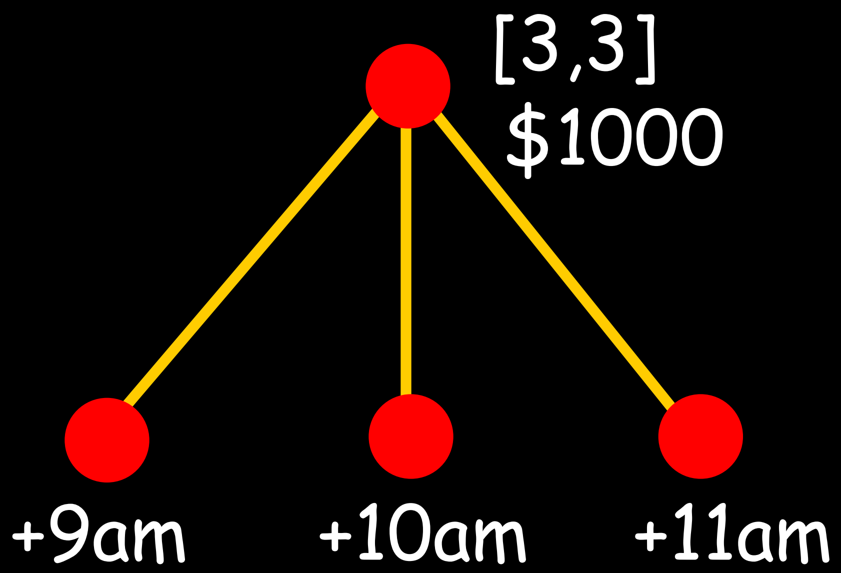

Example 2: "xor"

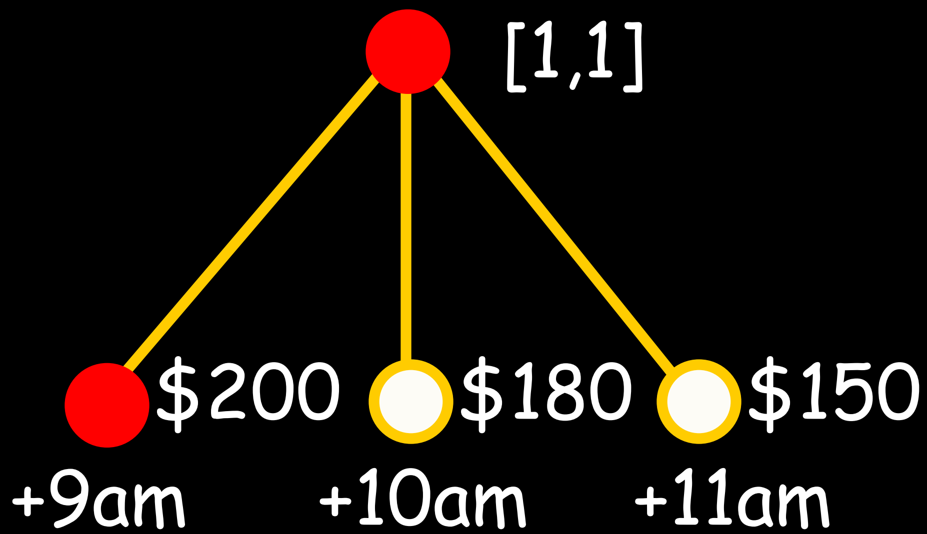

Example 3: “xor of and”

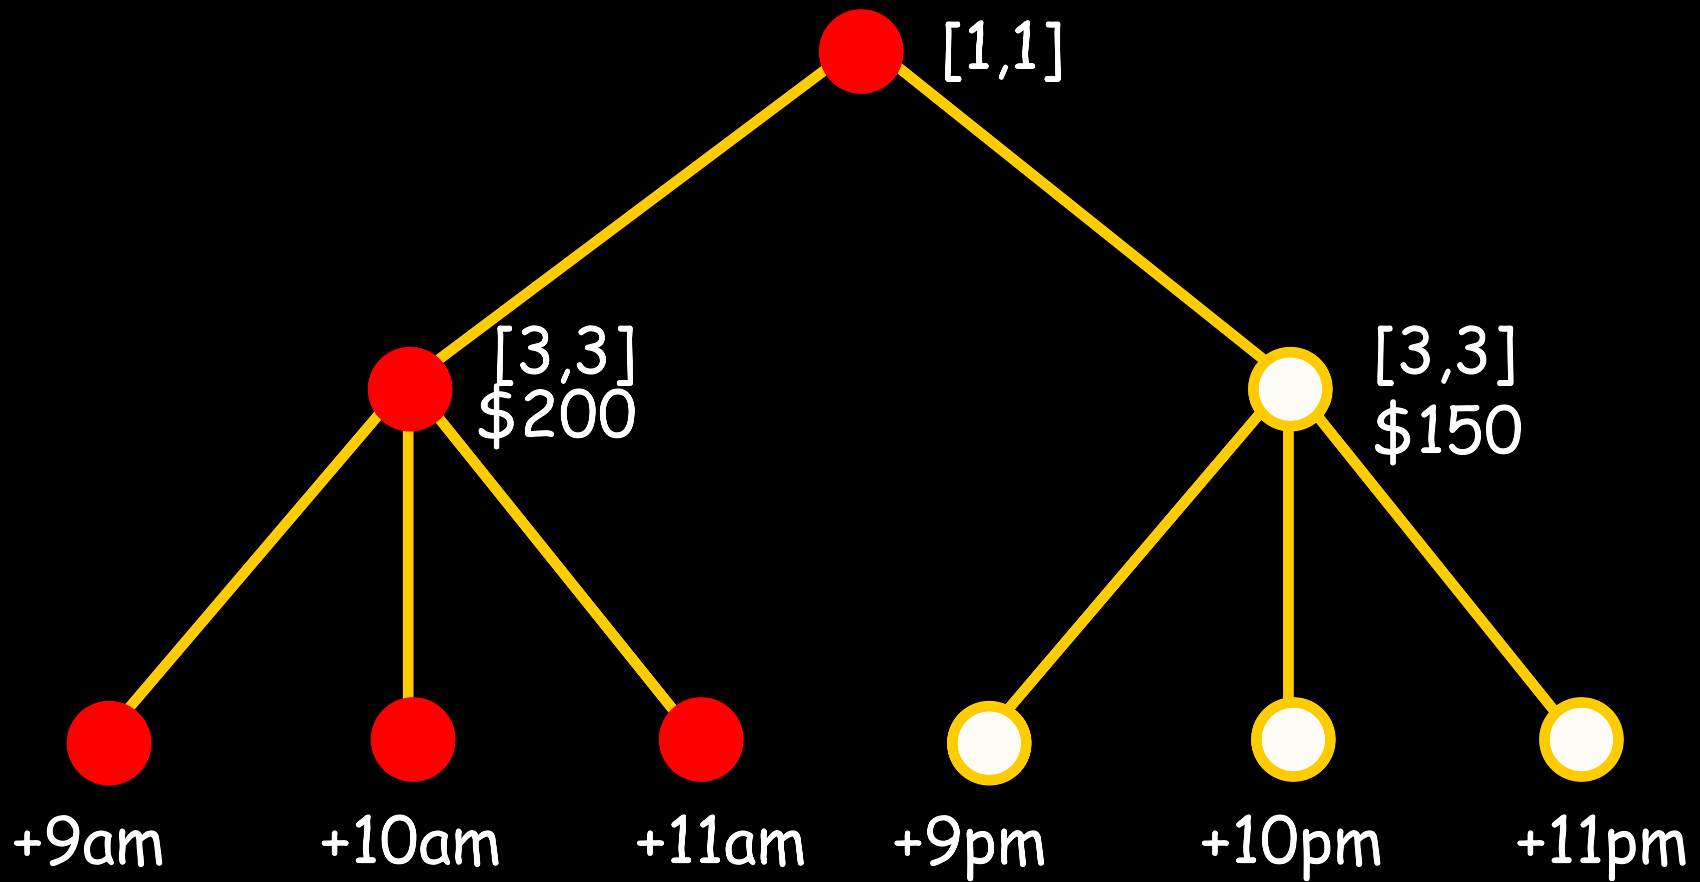

Example 4: “choose”

- `IC[x,y]`: accept an allocation in which *at least x* and *at most y* children are “satisfied”

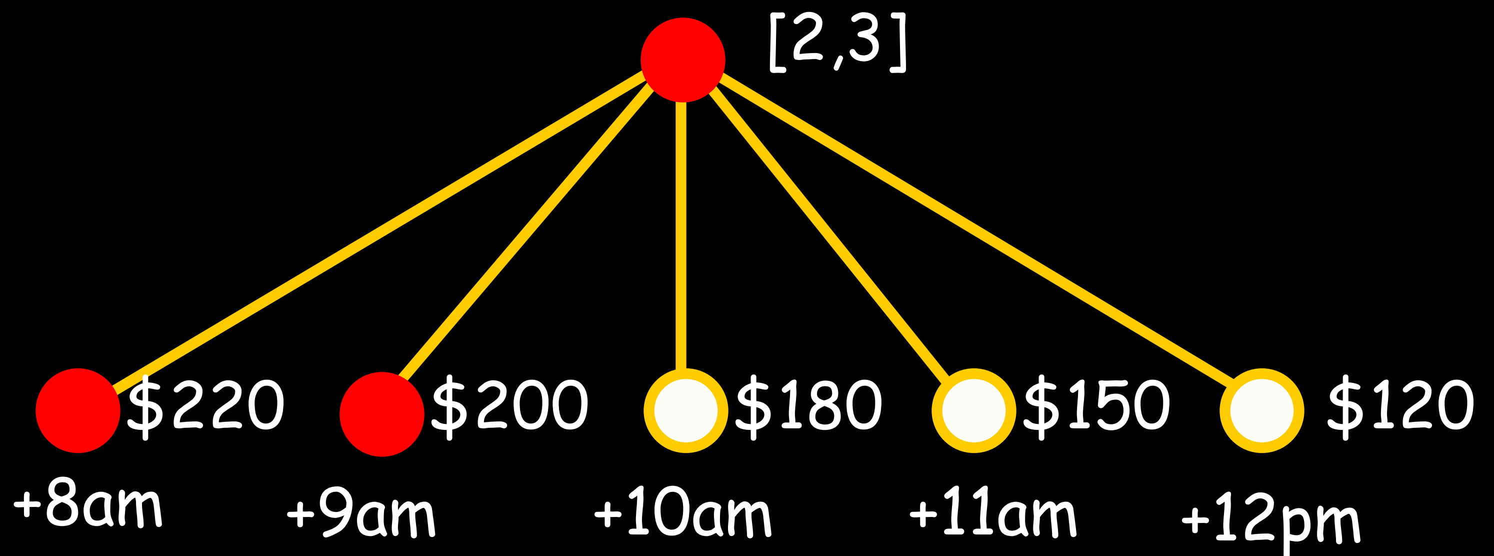

Example 5: “swap”

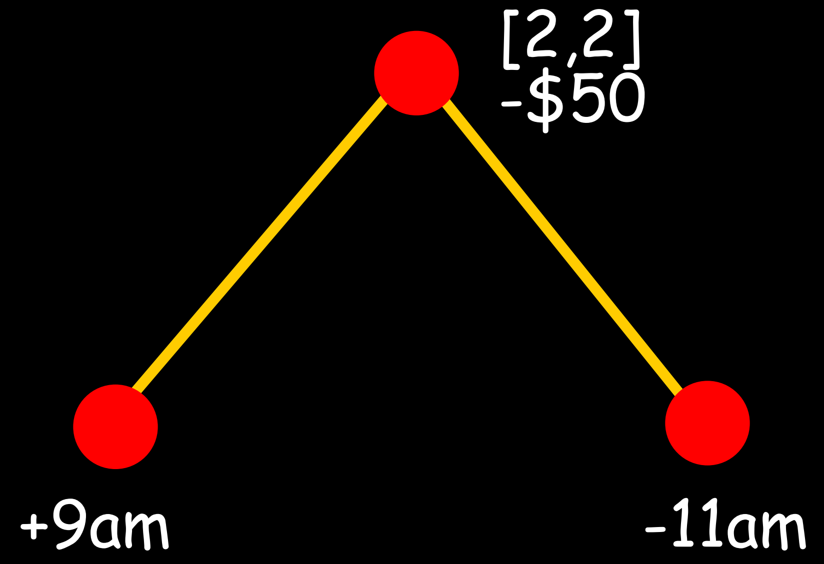

Example 6: “contingent sale”

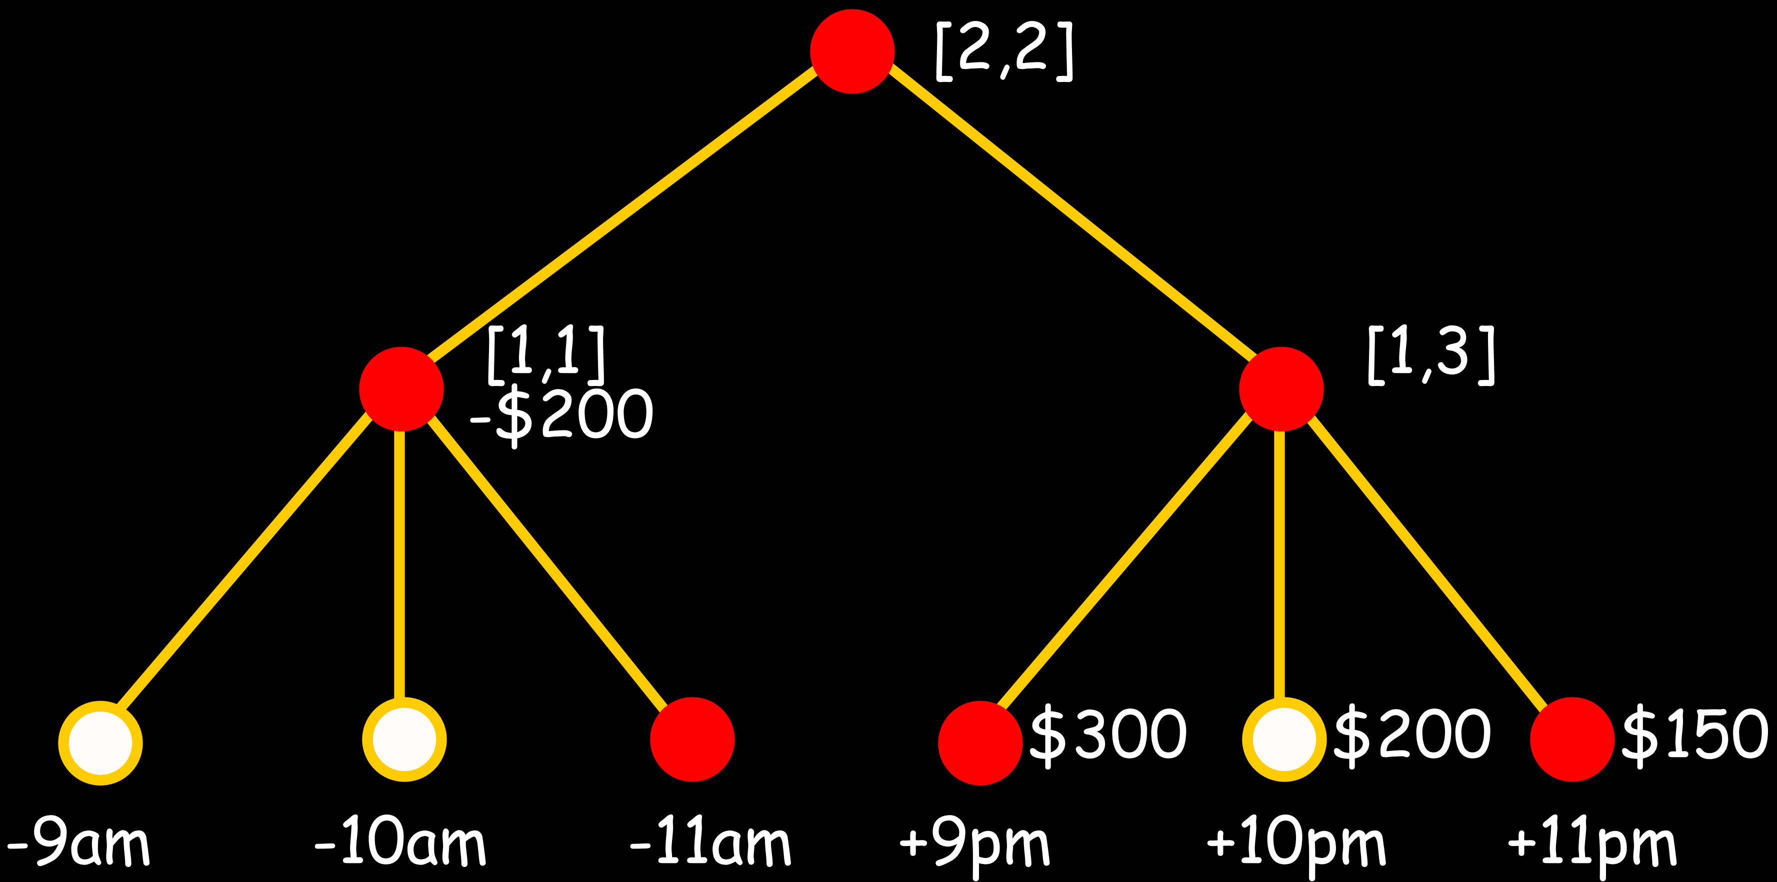

Winner Determination

- Goods: $\{1,\ldots,m\}$. Agents: $\{1,\ldots,n\}$
- Trades: $\lambda_{ij} \in Z$ Initial allocation: $x^0_{ij} \in Z$
- Winner determination:

$$\max \sum_i v_i(\lambda_i)$$

$$\text{s.t.} \quad \left.\begin{array}{l} \lambda_{ij} + x^0_{ij} \geq 0, \ \forall i \ \forall j \\ \sum_i \lambda_{ij} \leq 0, \ \forall j \\ \lambda_{ij} \in Z \end{array}\right\} \lambda \in \mathrm{feas}(x^0)$$

Value given λ_i?

$$\max_{sat_i\{\beta\}} \sum_{\beta\in T} v_i(\beta) sat_i(\beta)$$

$$\text{s.t. } \sum_{\beta\in Leaf(i)} q_{ij}(\beta) sat_i(\beta) \le \lambda_{ij},\ \forall j \qquad (3)$$

$$IC_{x,i}(\beta) sat_i(\beta) \le \sum_{\beta'\in child(\beta)} sat_i(\beta') \le IC_{y,i}(\beta) sat_i(\beta),\ \forall \beta \notin Leaf(i) \qquad (4)$$

$$sat_i(\beta) \in \{0,1\}$$

Concise WD Formulation

$\max_{\lambda,\mathrm{sat}} \sum_i \sum_{\beta \in T_i} v_i(\beta)\mathrm{sat}_i(\beta)$

s.t. (feas), (TBBL semantics)

Linear in size of TBBL trees

ICE: Proxied Exchange

(Parkes et al. 2007)

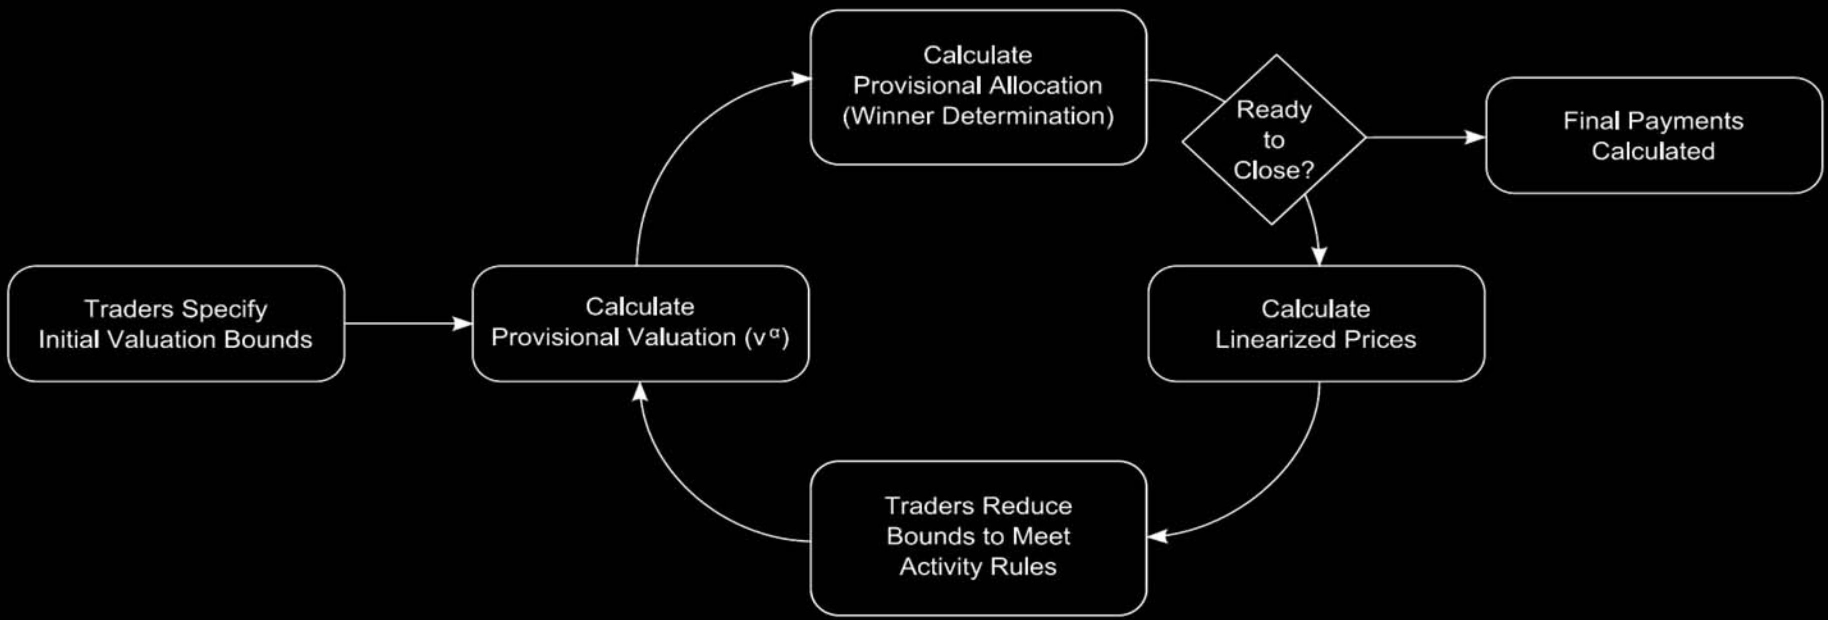

Activity Rule

- Show one trade is weakly better then all others

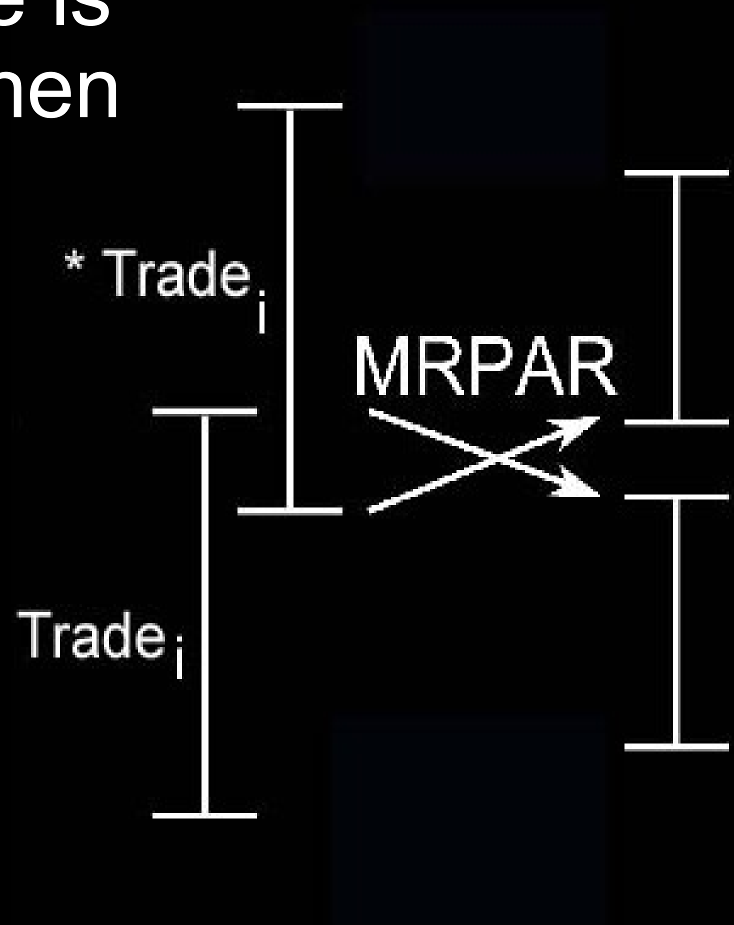

- + another activity rule in later stages

Scalability (I)

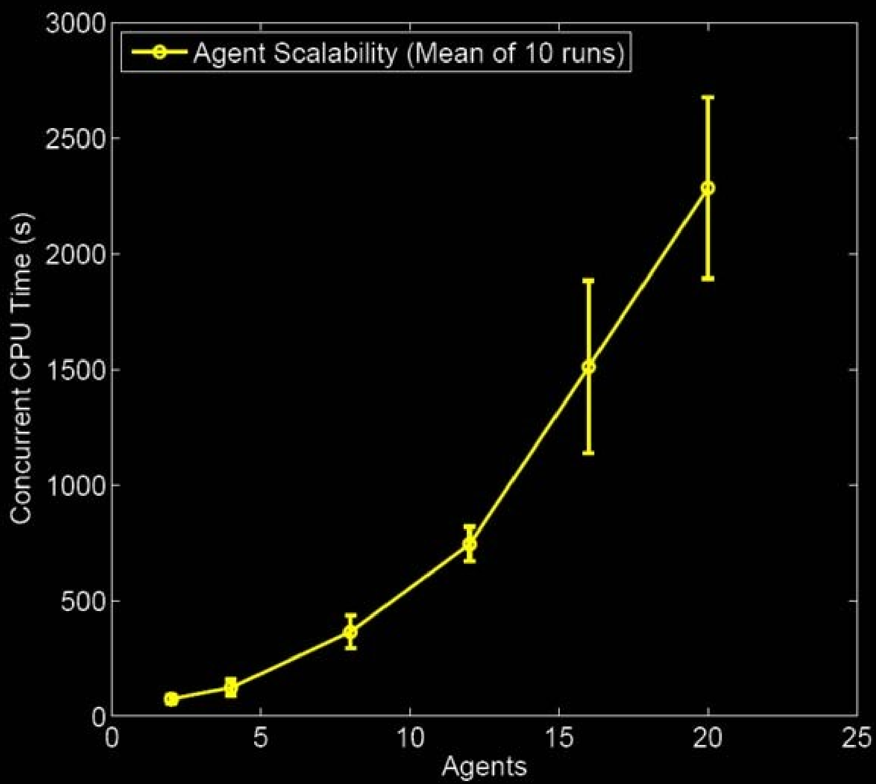

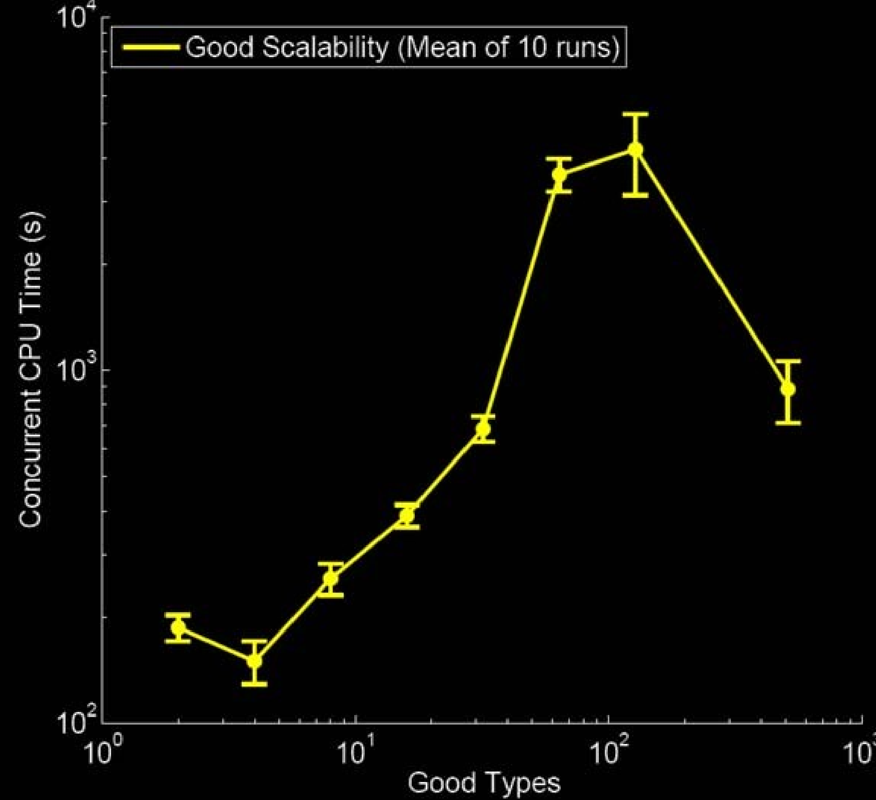

Scalability (II)

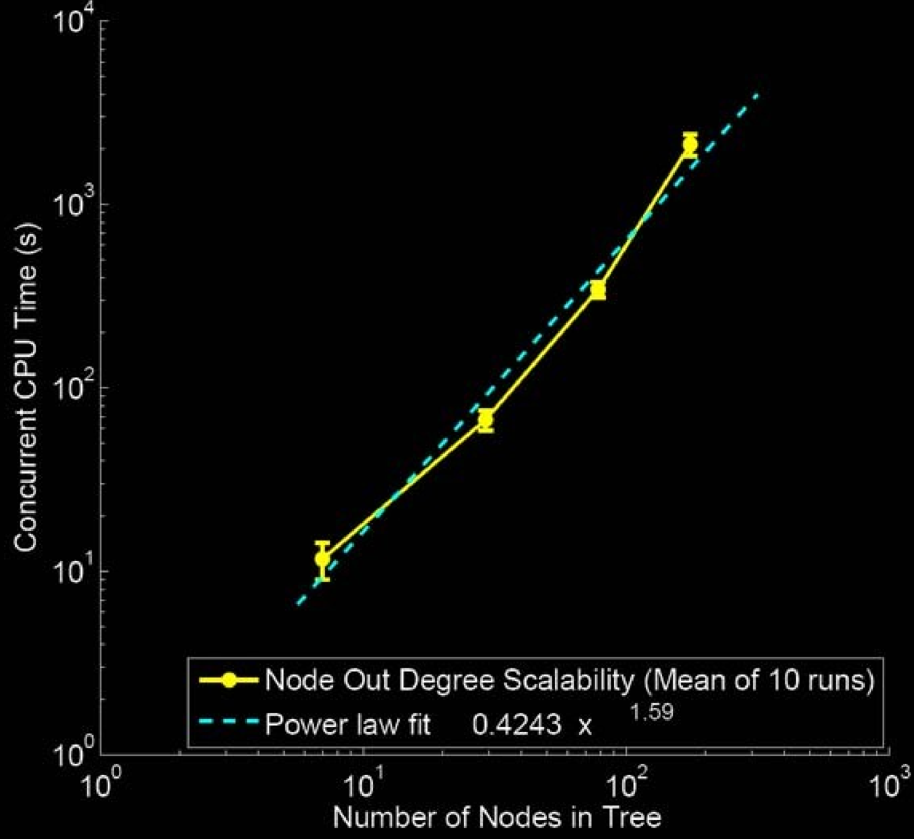

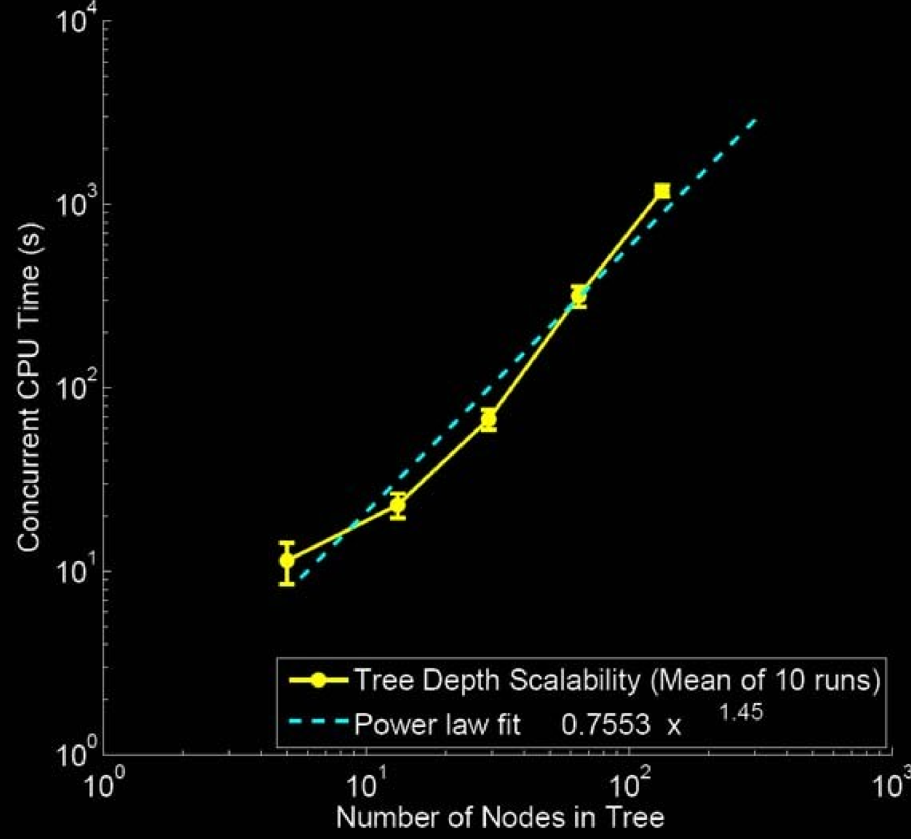

Price Feedback

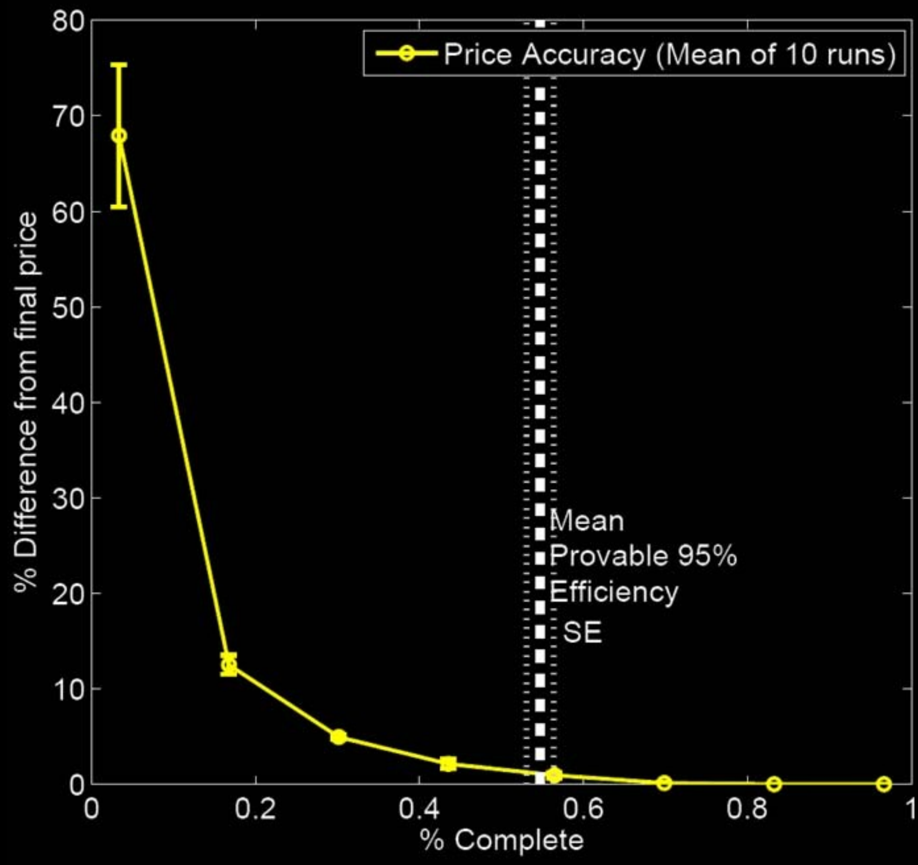

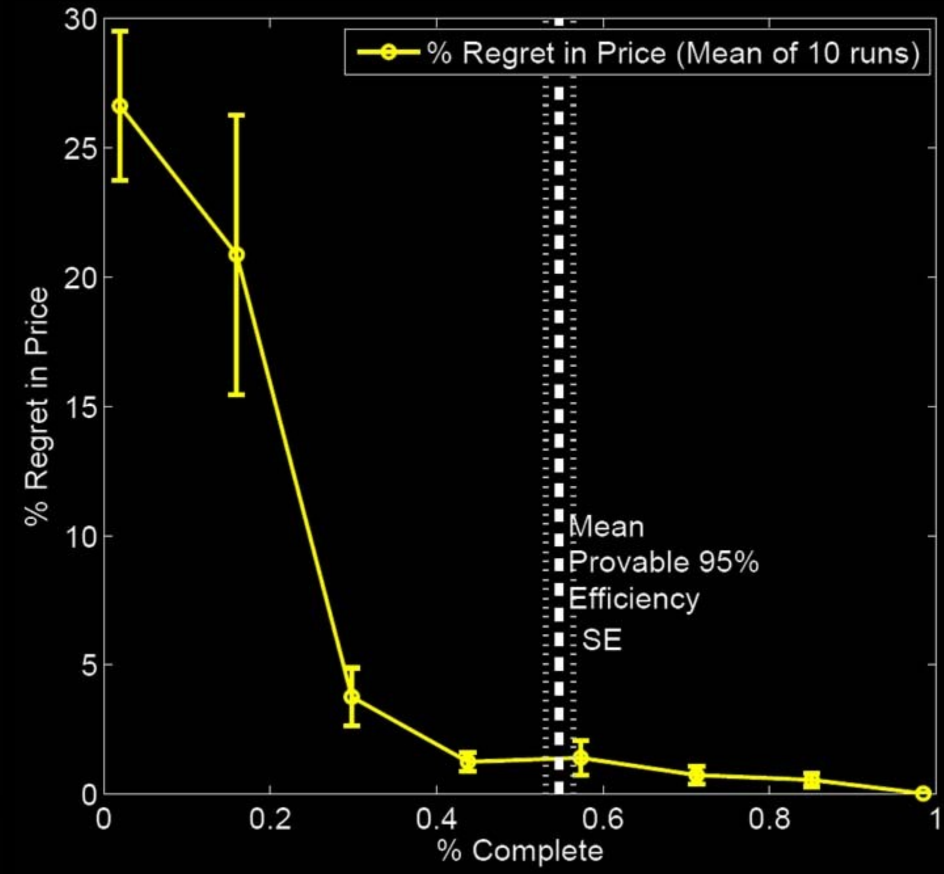

Payments?

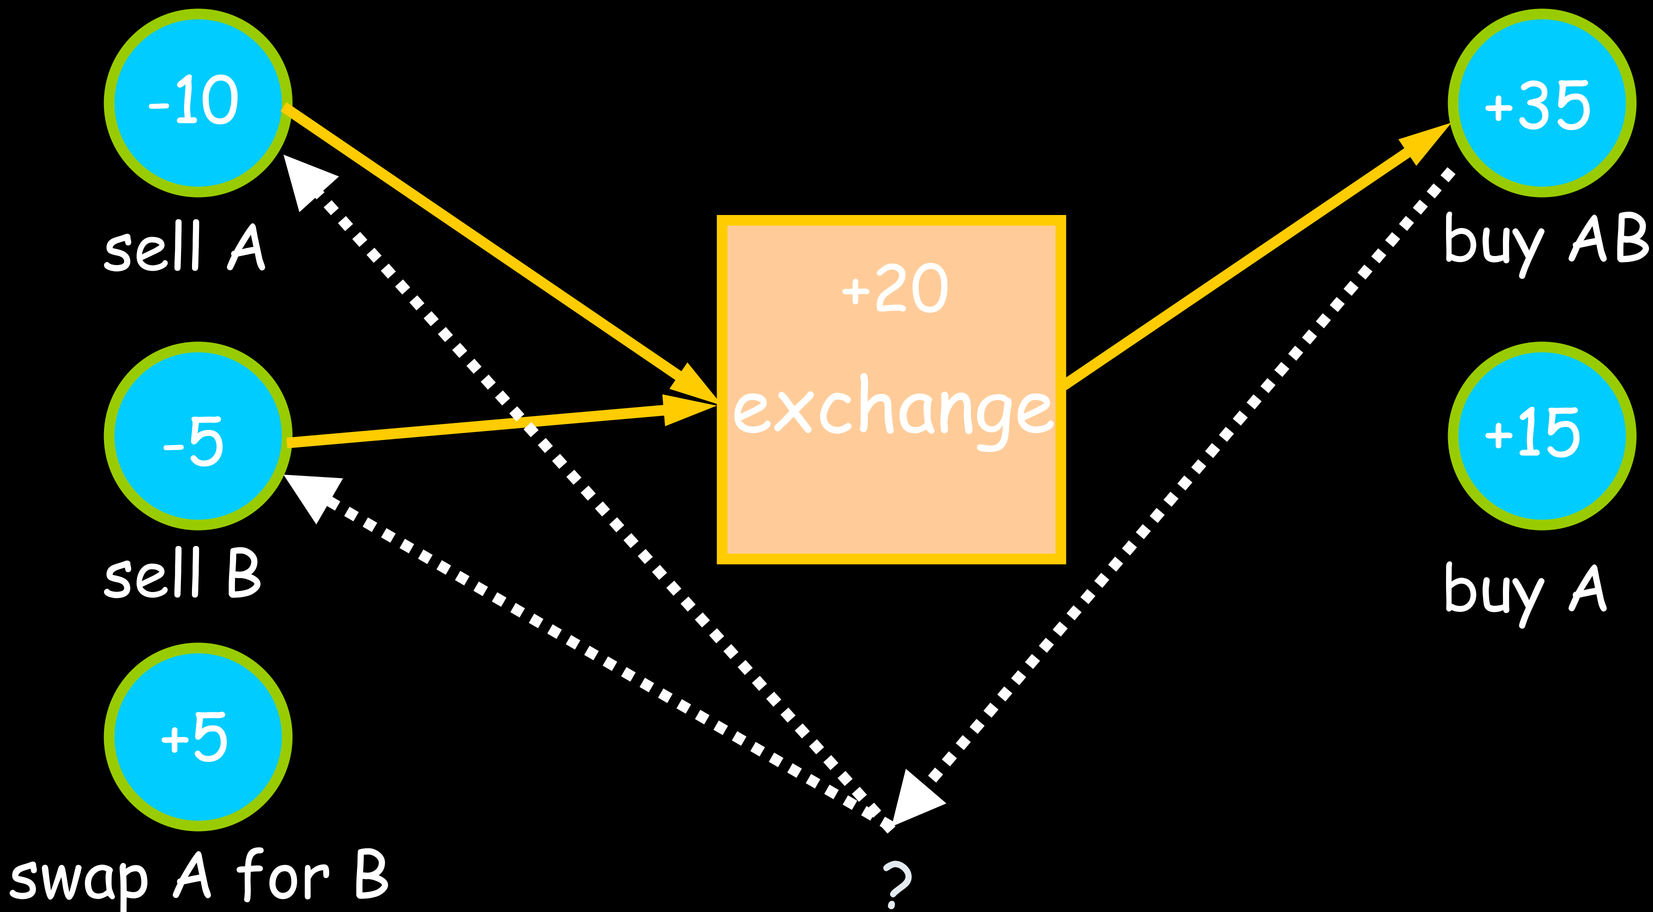

Payments: VCG?

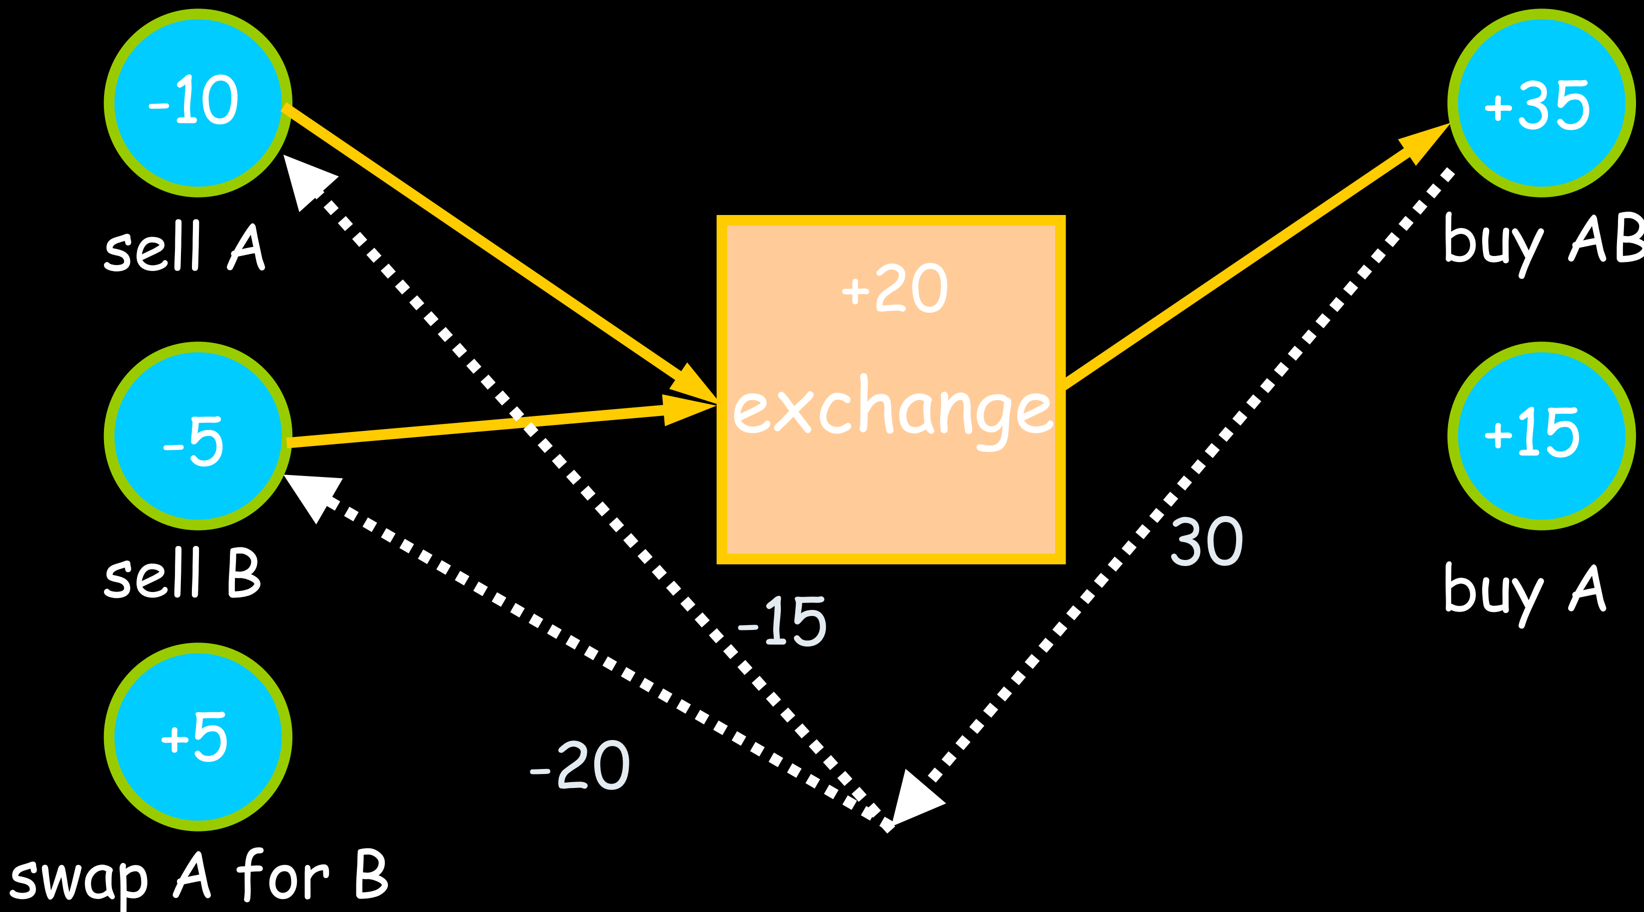

- Myerson-Satterthwaite impossibility
  - EFF, No-deficit, IR and BNIC

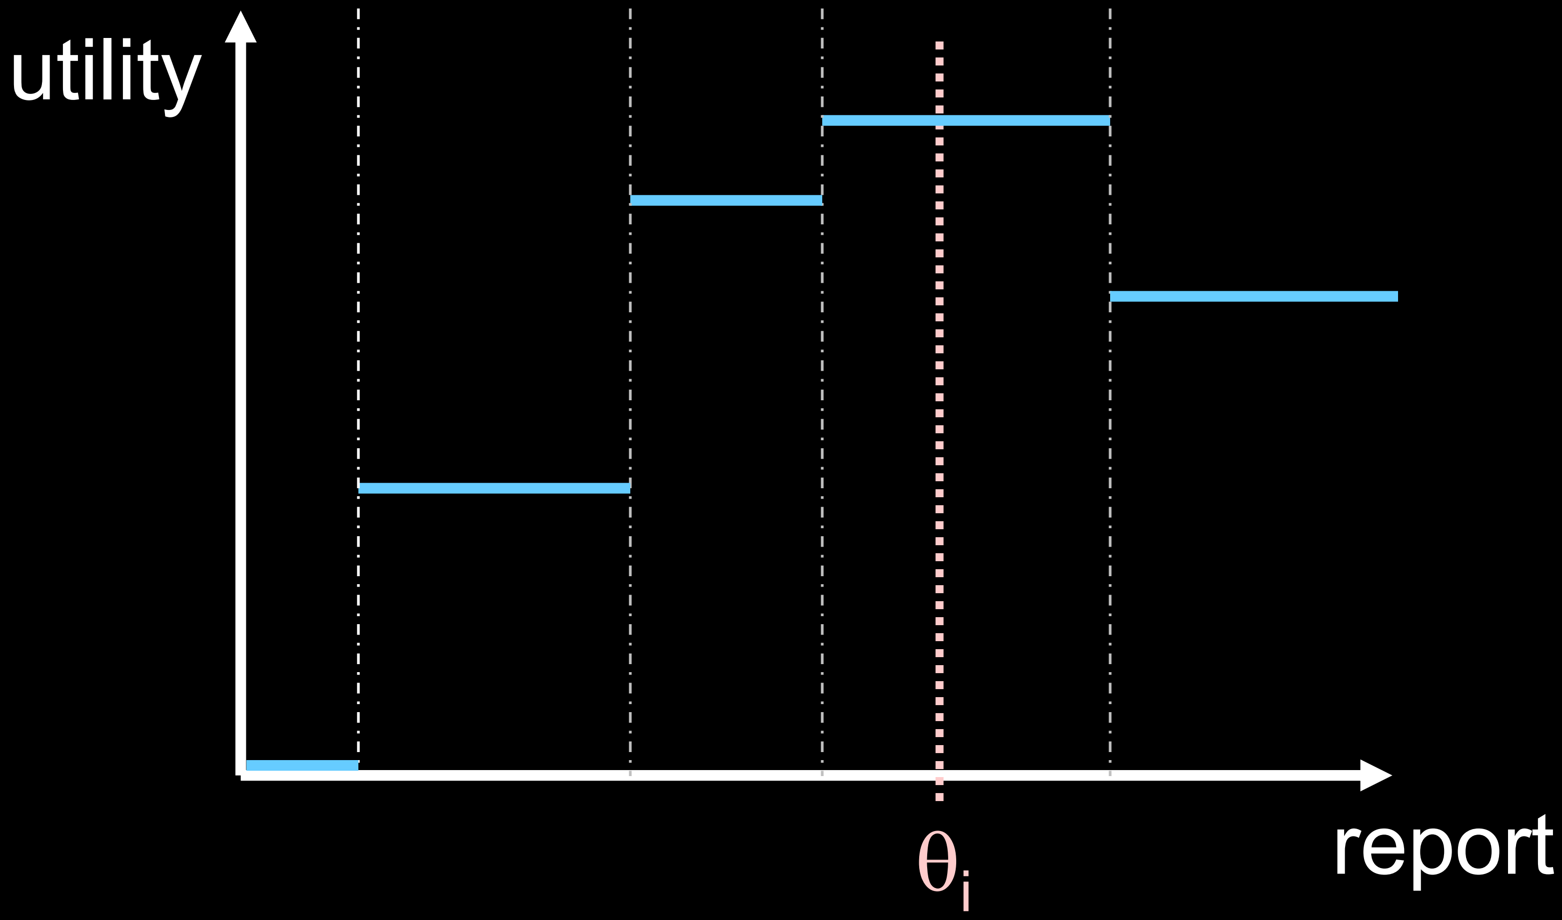
utility
report
θi

Approx SP: GSP

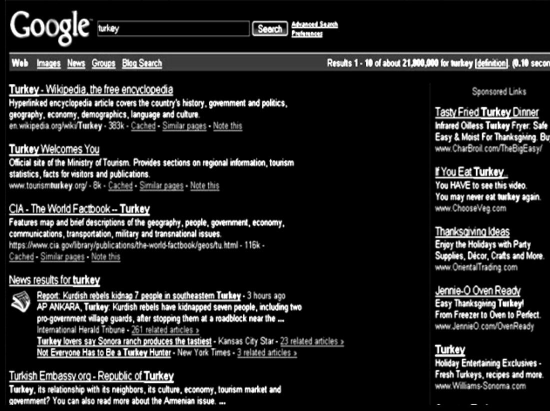

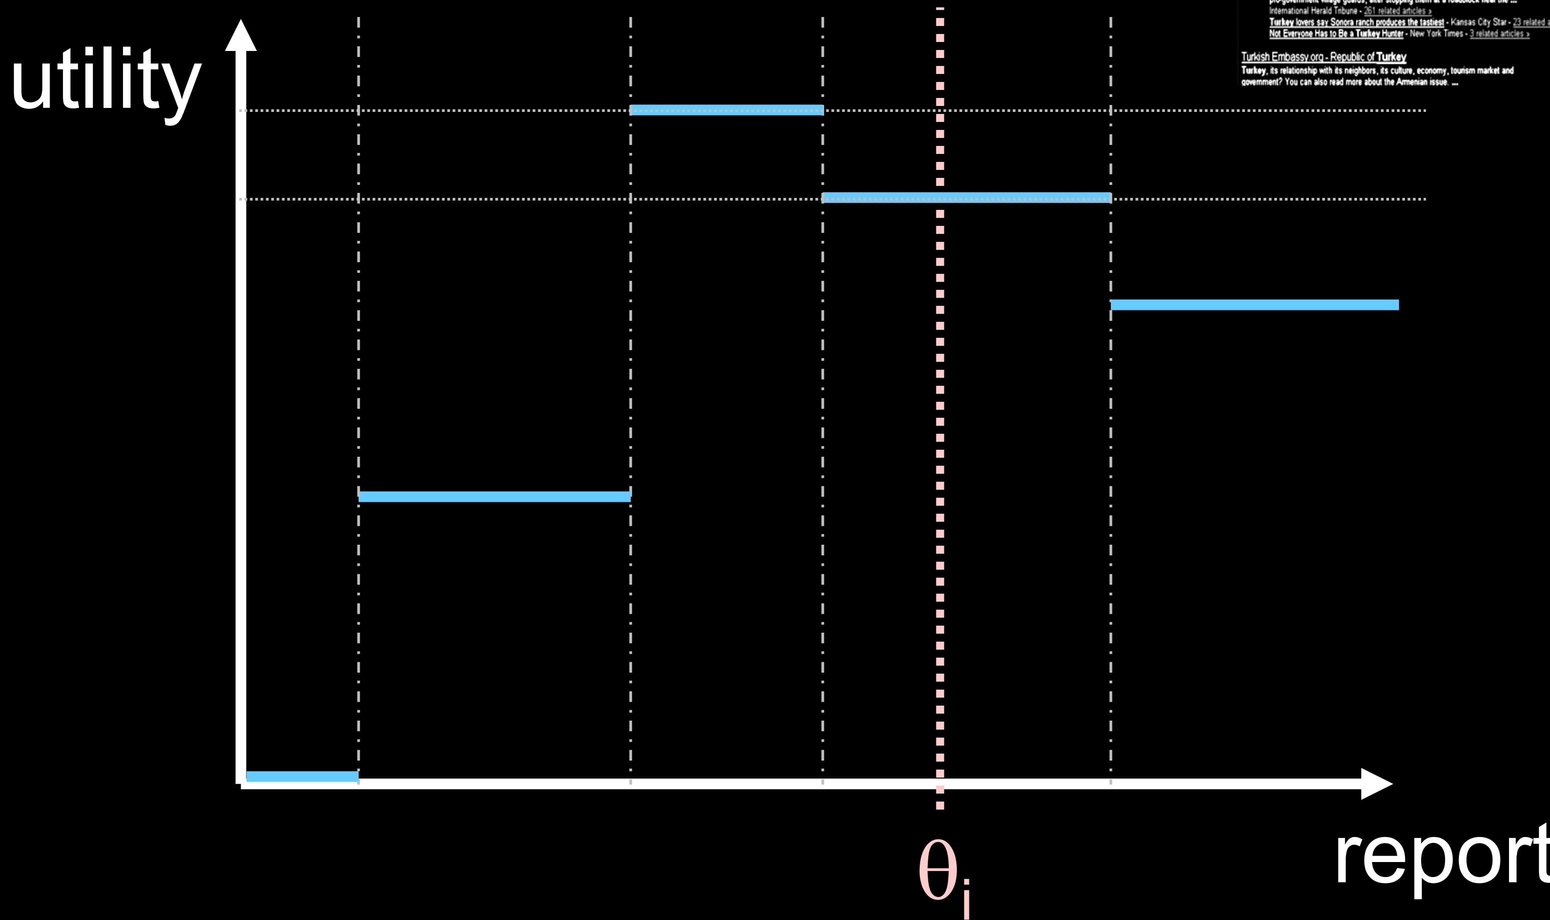

How should we set payments in a CE that clears straightforwardly based on bids?

($\equiv$ how should we make the mechanism “maximally incentive compatible”?)

Relaxing away from SP…

- We like SP for reasons of
  - equity (Roth’03, Pathak and Sonmez’08)
  - *simplify reasoning*
  - can predict properties of the mechanism

Relaxing away from SP…

- We like SP for reasons of
  - equity (Roth’03, Pathak and Sonmez’08)
  - *simplify reasoning*
  - can predict properties of the mechanism
- But it is generally hard to obtain
- And, can be provably bad along other dimensions ☹
    - e.g., CAs with complements
      (Ausubel & Milgrom’06, Rastegeri, Condon, & Leyton-Brown’10)

Example: Course Allocation
(Budish and Cantillon'08)

- Random Serial Dictatorship
  - basically unique amongst SP mechanisms (Papai'01)
- HBS mechanism:
  - snake back and forth, pick one at a time

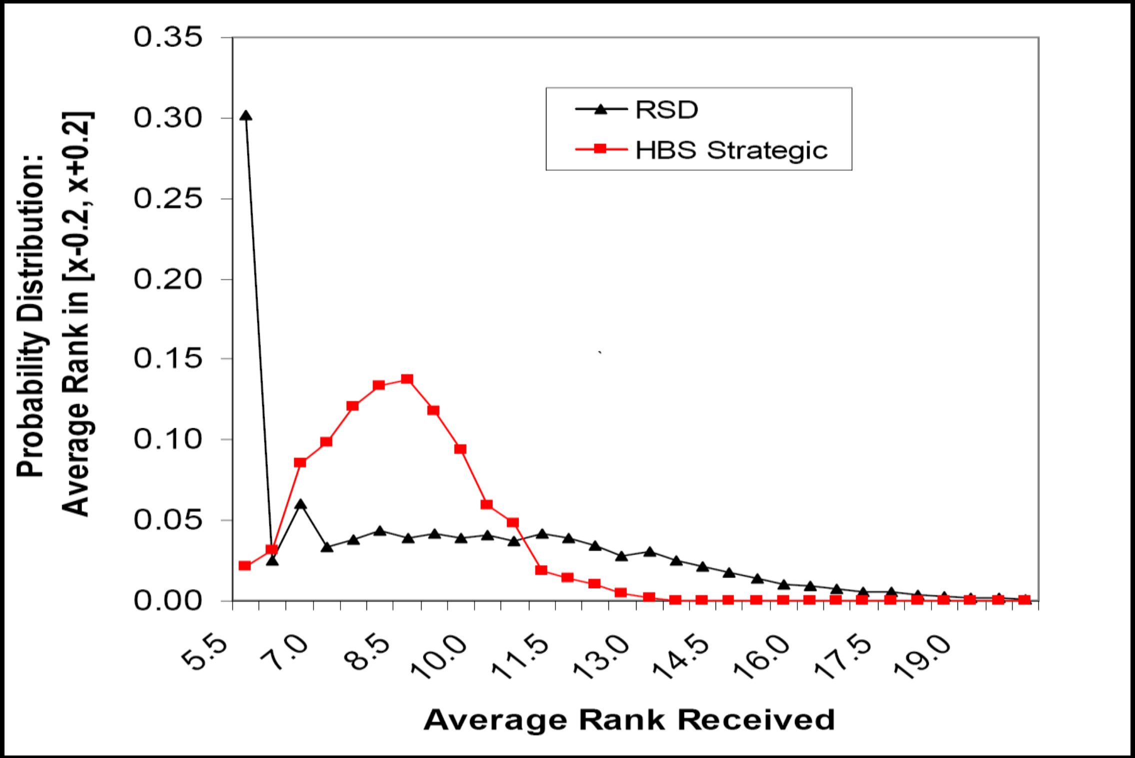

"callousness of RSD"

(allocating 10 courses)

Old Favorite: Min Max Regret

- Regret = best utility – actual utility
- Maximally SP: minimizes max regret across agents on every instance
- ϵ-SP: max regret $\leq \epsilon$

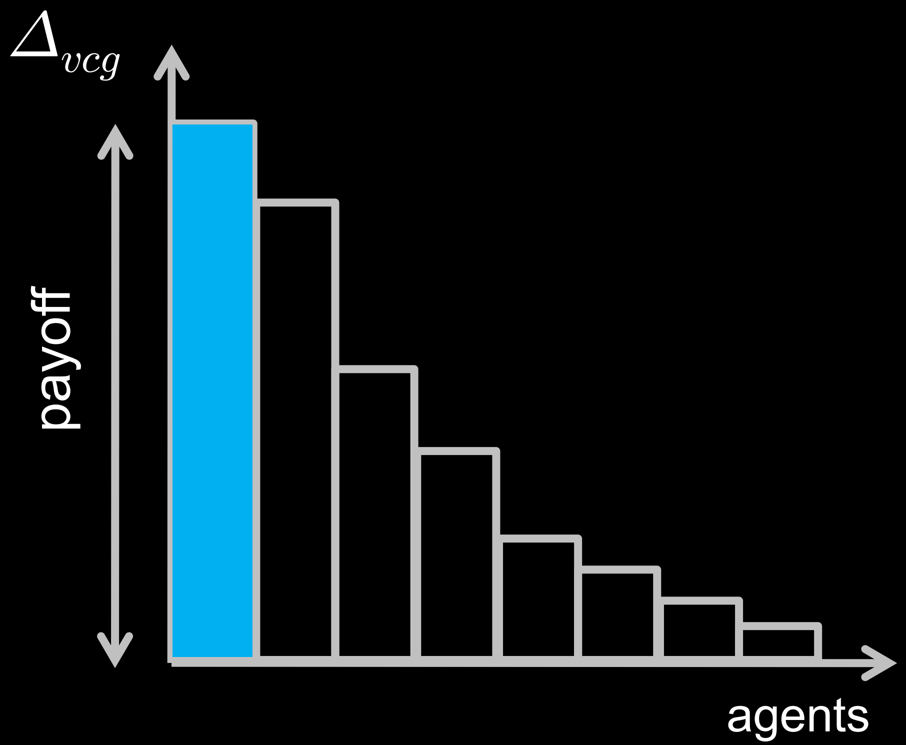

Δ_{vcg}
payoff
agents

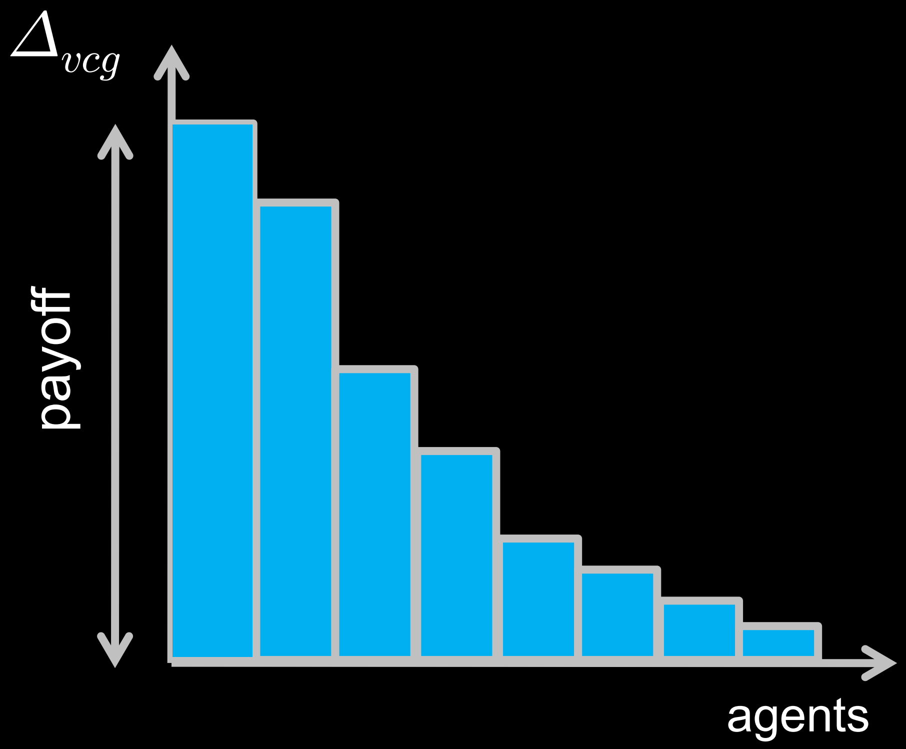
Δ_{vcg}
payoff
agents

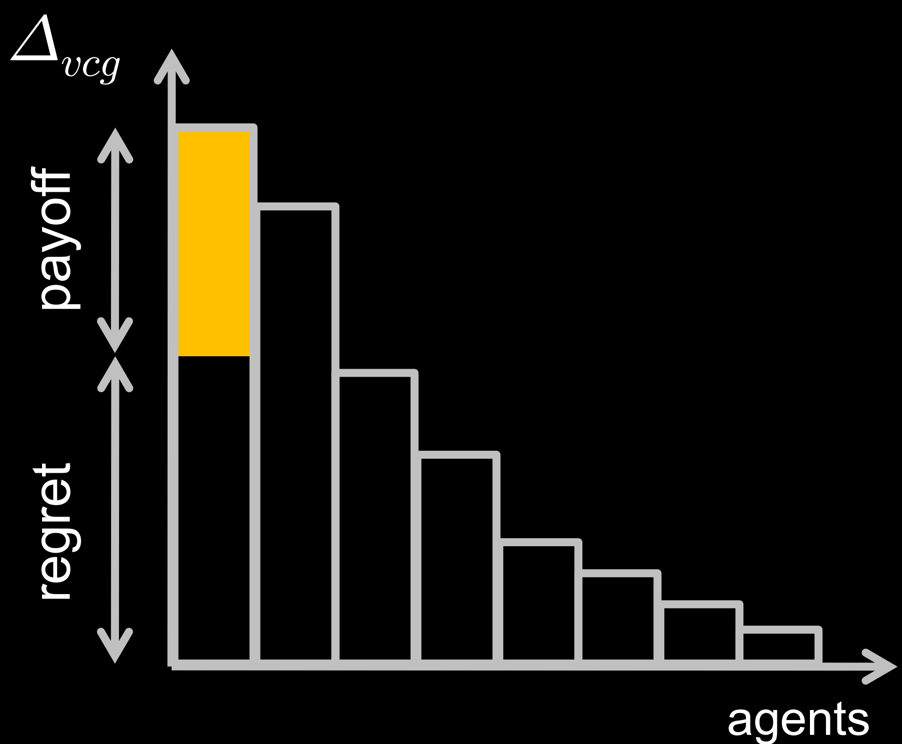

Δ_{vcg}
payoff
regret
agents

Two mechanism rules

(Parkes, Kalagnanam and Eso '01)

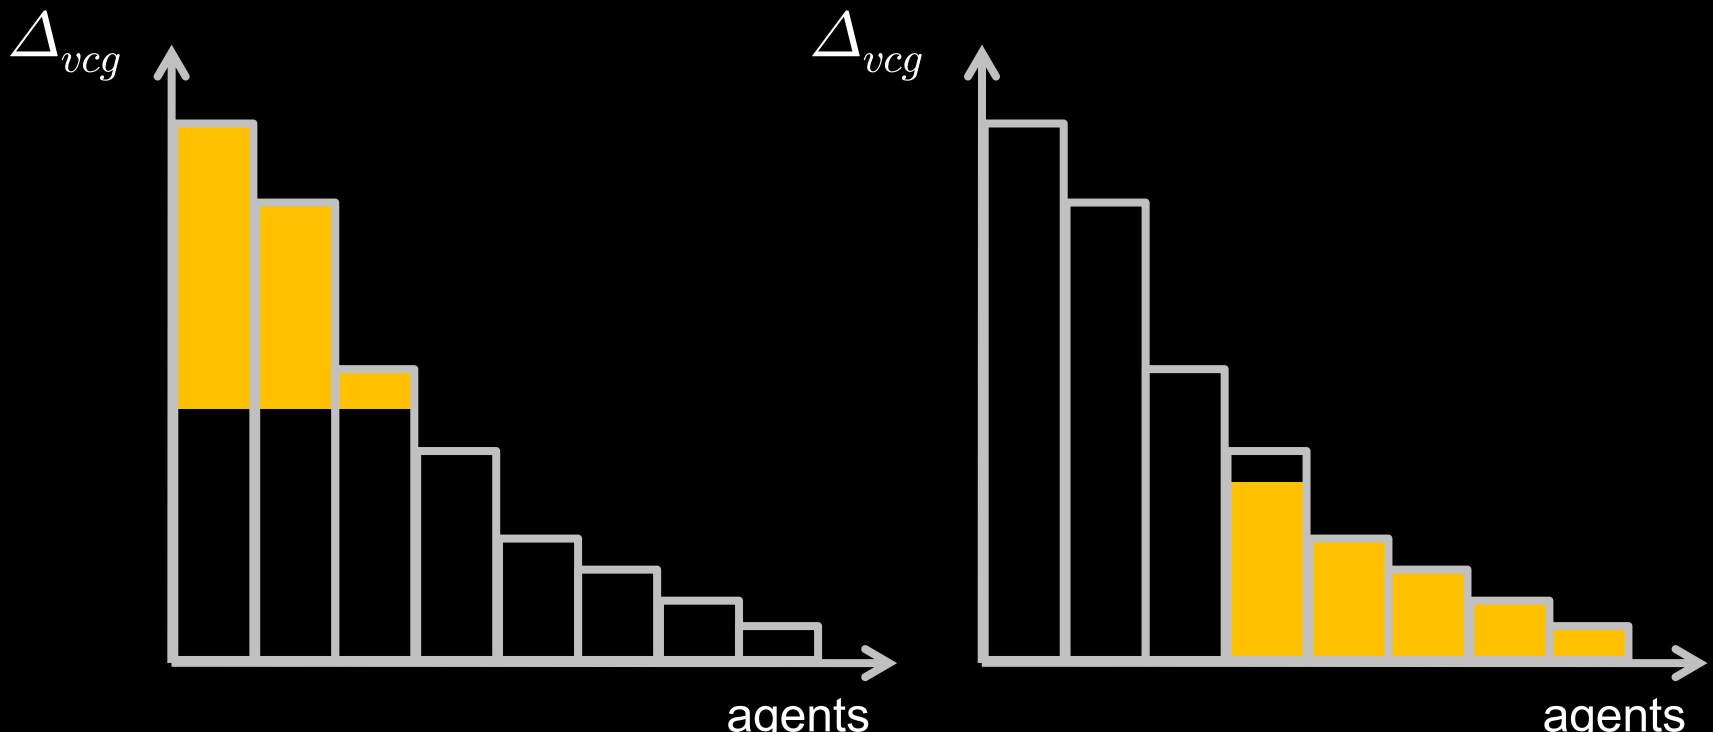

Threshold rule
(min max regret)

Small rule
max #(regret=0)

Back to Example

$\Delta_{vcg} = (5,15,5)$ Surplus 20

- *Threshold*: $\Delta_1 = 3.33$, $\Delta_2 = 13.33$, $\Delta_3 = 3.33$
  - payments (-13.33, -18.33, +31.67)
  - regret = 1.33 for all agents
- *Small*: $\Delta_1 = 5$, $\Delta_2 = 10$, $\Delta_3 = 5$
  - payments (-15, -15, +30)
  - regret = 0, 5, 0

- **Theorem**. Threshold rule minimizes ex post opportunity for gain across simple CEs
  - "truthful most often" (assume cost C_d) [Milgrom]

Compute approx BNE

(Related: Vorobeychik et al.)

- Single-minded CEs
- Need a way to compute approximate, restricted BNE
- Approach: assume piecewise linear, symmetric strategy profiles

- Buy: bid (1+α) v
- Sell: bid (1+α) v
  - Use (α_1, α_2, α_3)

Approximate BNE Analysis

(Lubin & Parkes ’09)

	strategy			efficiency		
Rule	**Dec.**	**Uni.**	**Sup.**	**Dec.**	**Uni.**	**Sup.**
VCG	*0.0*	*0.0*	*0.0*	*100*	*100*	*100*
Two Triangle	**0.1**	**0.4**	**5.6**	**99.99**	**100**	**97.95**
Threshold	14.6	27.2	11.2	93.64	81.09	89.74
Small	**0.0**	**0.1**	**0.2**	**99.99**	**100**	**100**
No Discount	62.3	80.9	72.4	34.15	50.11	48.21

(For BNE, see Vorobeychik & Wellman’08, Rabinovich, Gerding, Polukarov & Jennings’09)

Distributional View: Payoffs

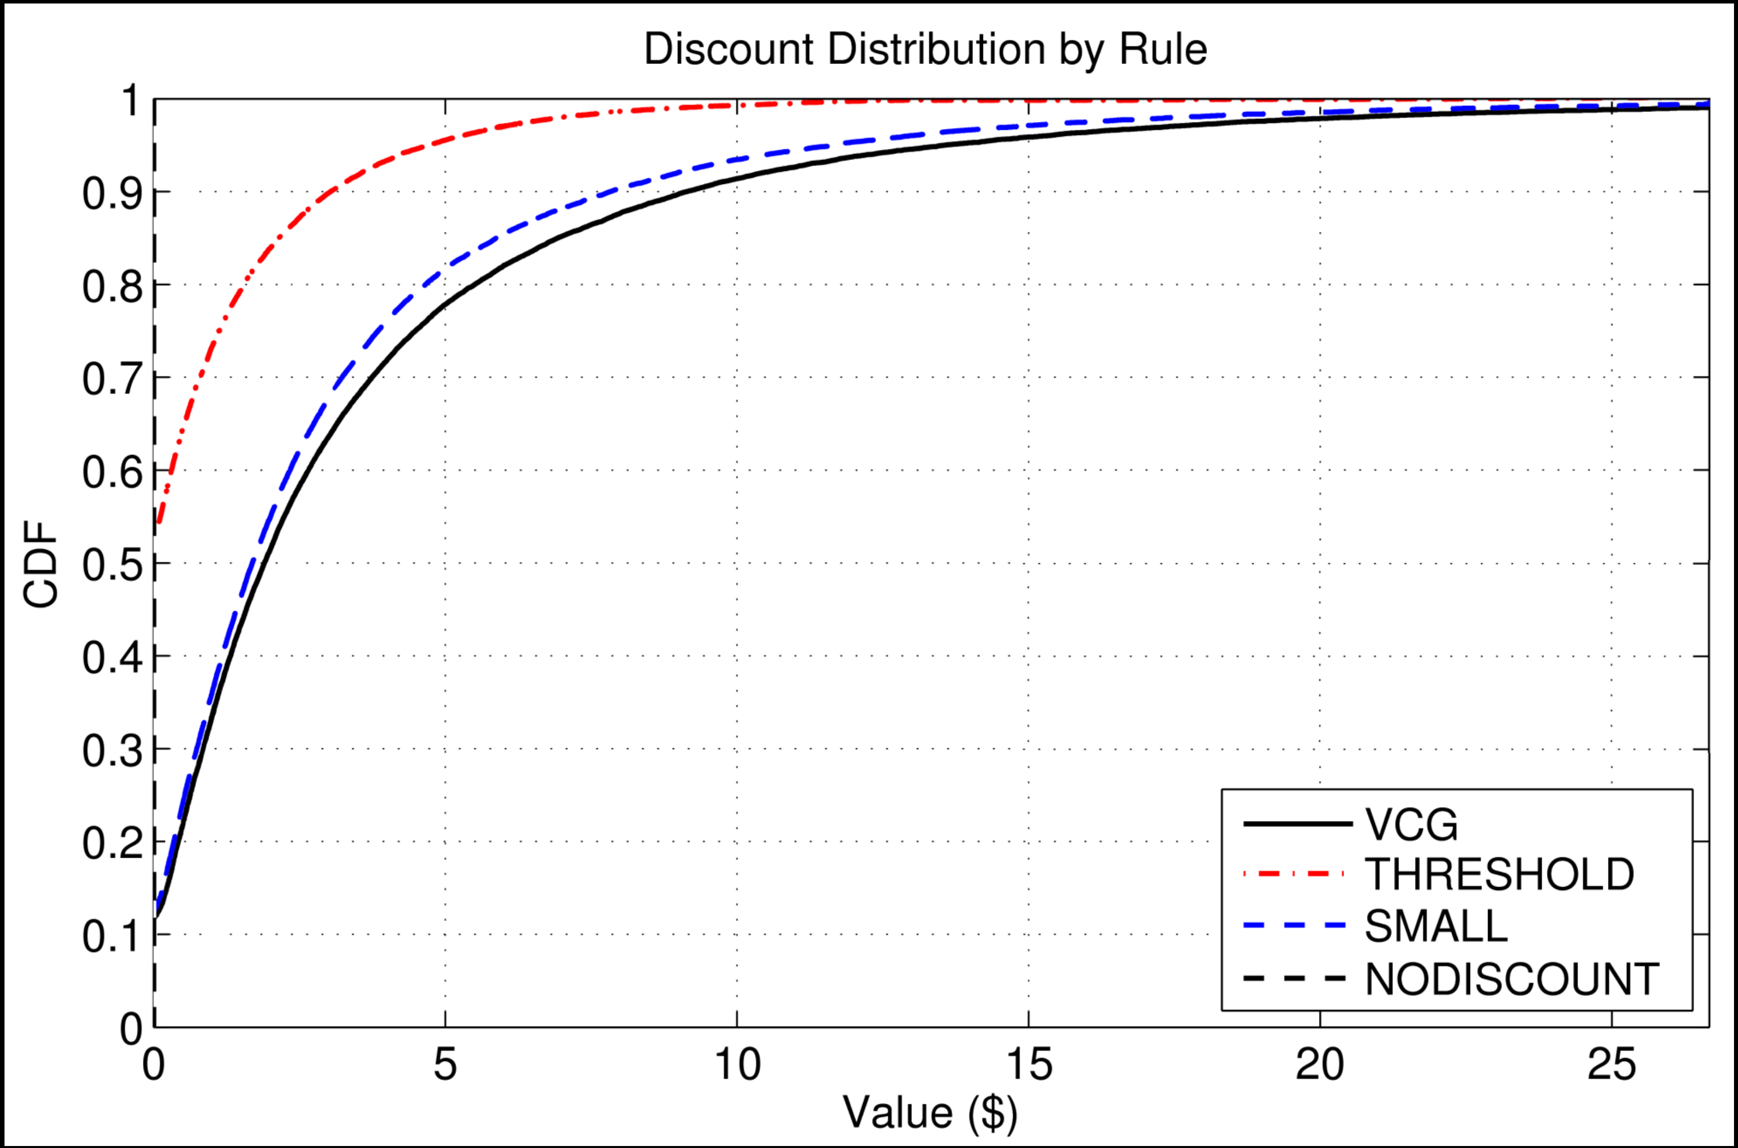

Regret Quantiles

(Lubin PhD ‘10)

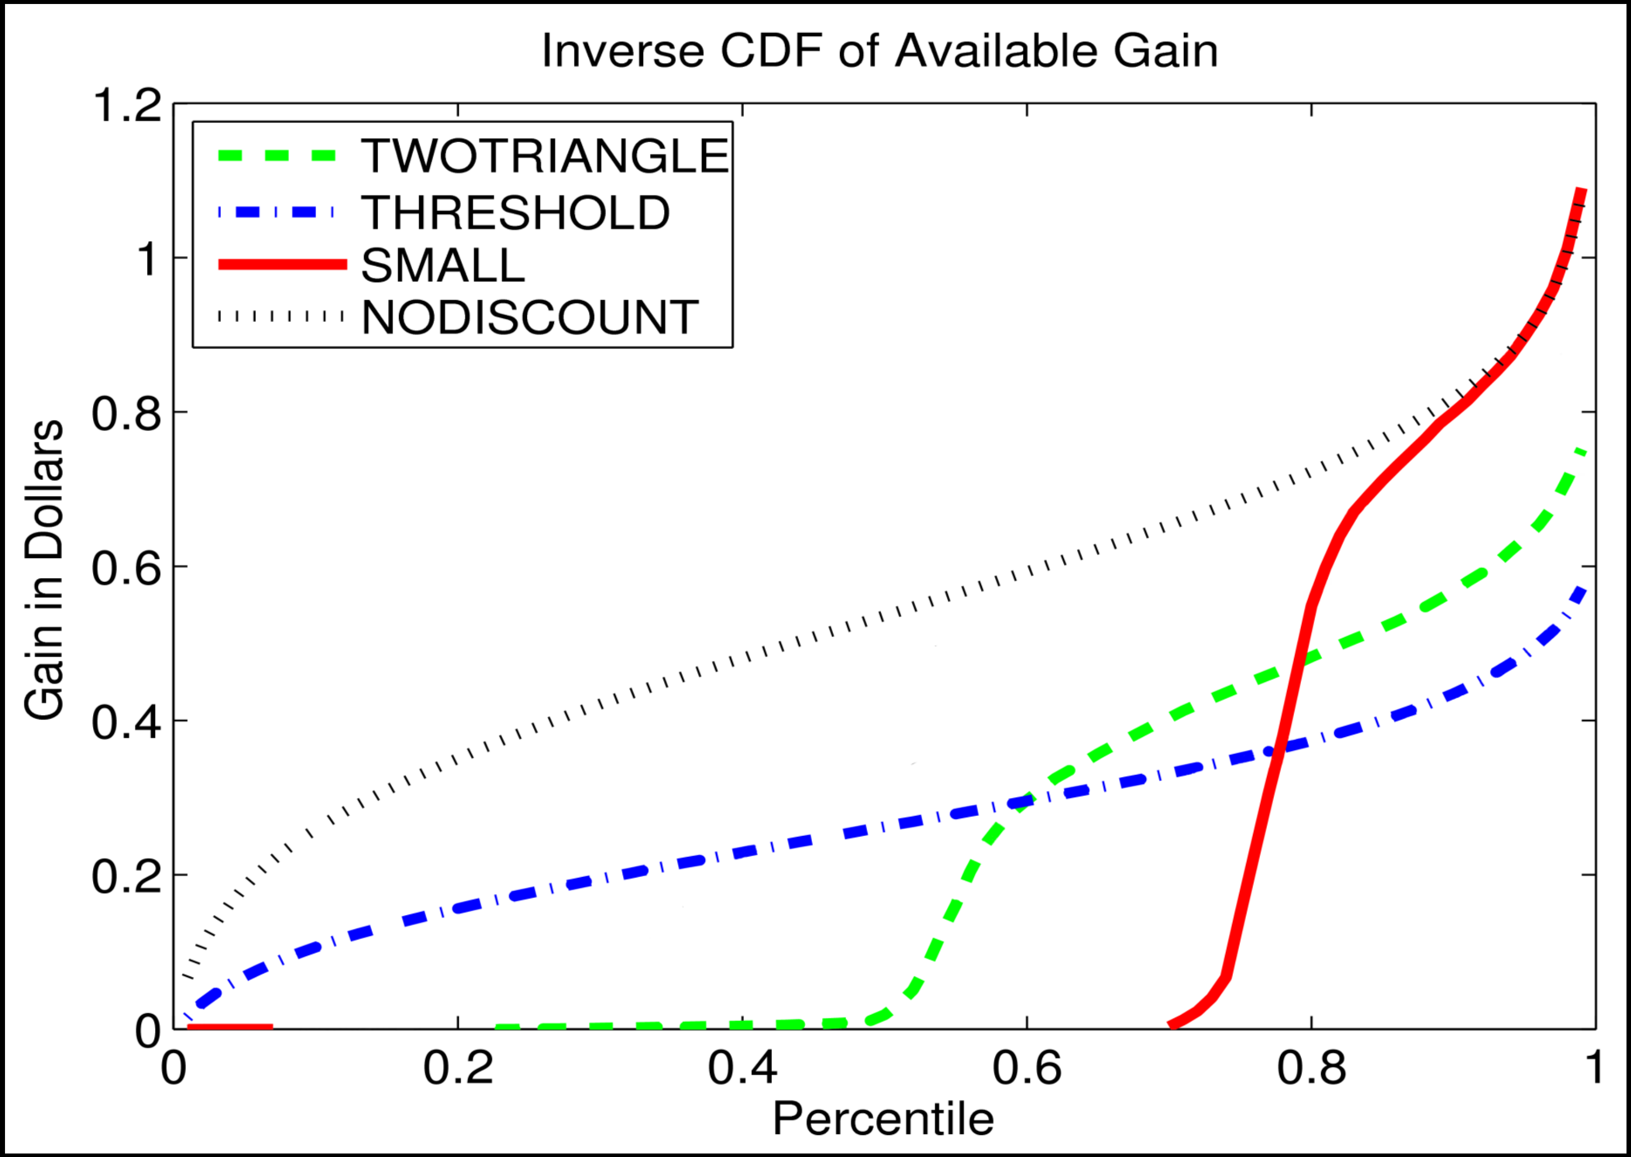

Regret Quantiles

(Lubin PhD ‘10)

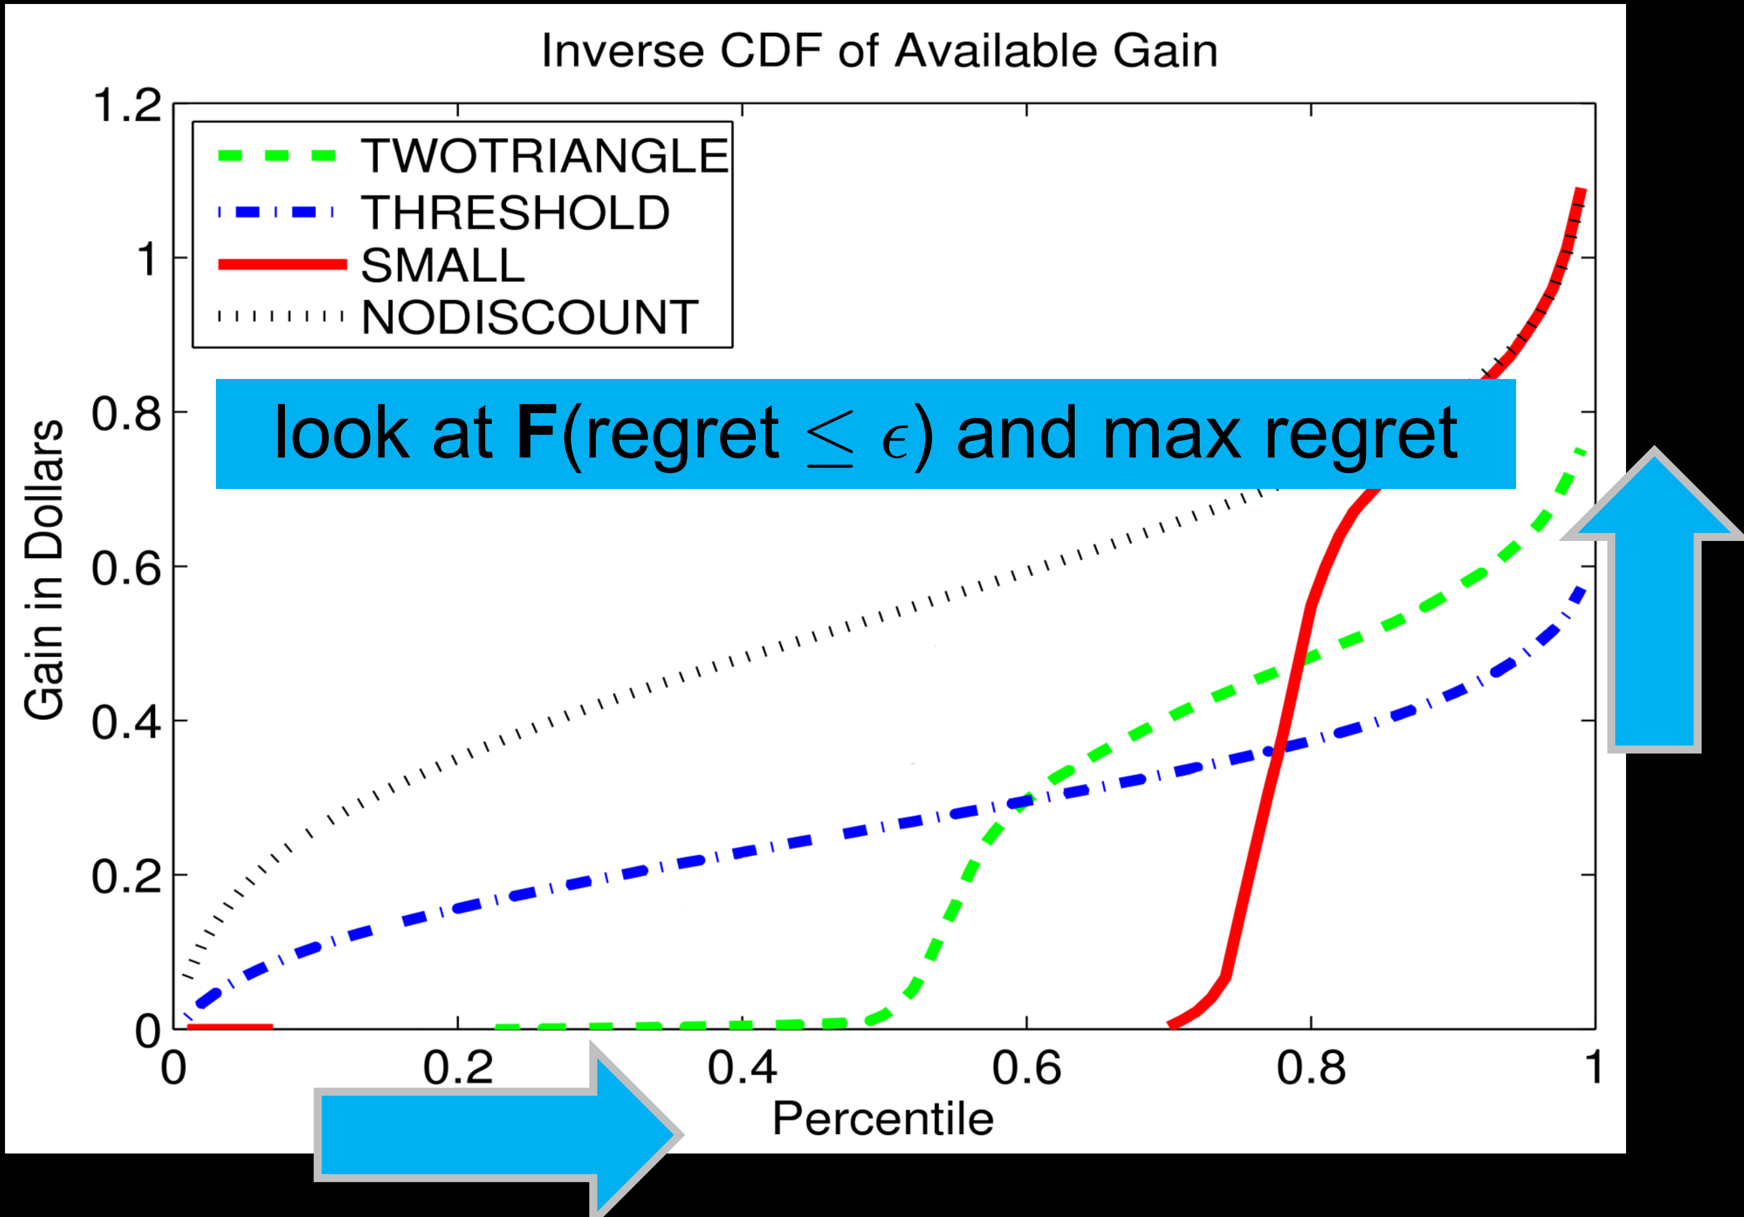

look at **F**(regret $\leq \epsilon$) and max regret

Hypothesis I

- Maximizing the number of agents with zero regret provides less *ex ante* incentive for strategic behavior than minimizing the maximum regret.

- (… proof would require reasoning about distributional properties)

Hypothesis II

- Consider a strategyproof “reference” mechanism M* with the same allocation rule but a different payment rule

- Reducing the divergence between the distribution on payoffs in M’ and the distribution on payoffs in M* reduces the *ex ante* incentive for strategic behavior.

Distributional View: Payoffs

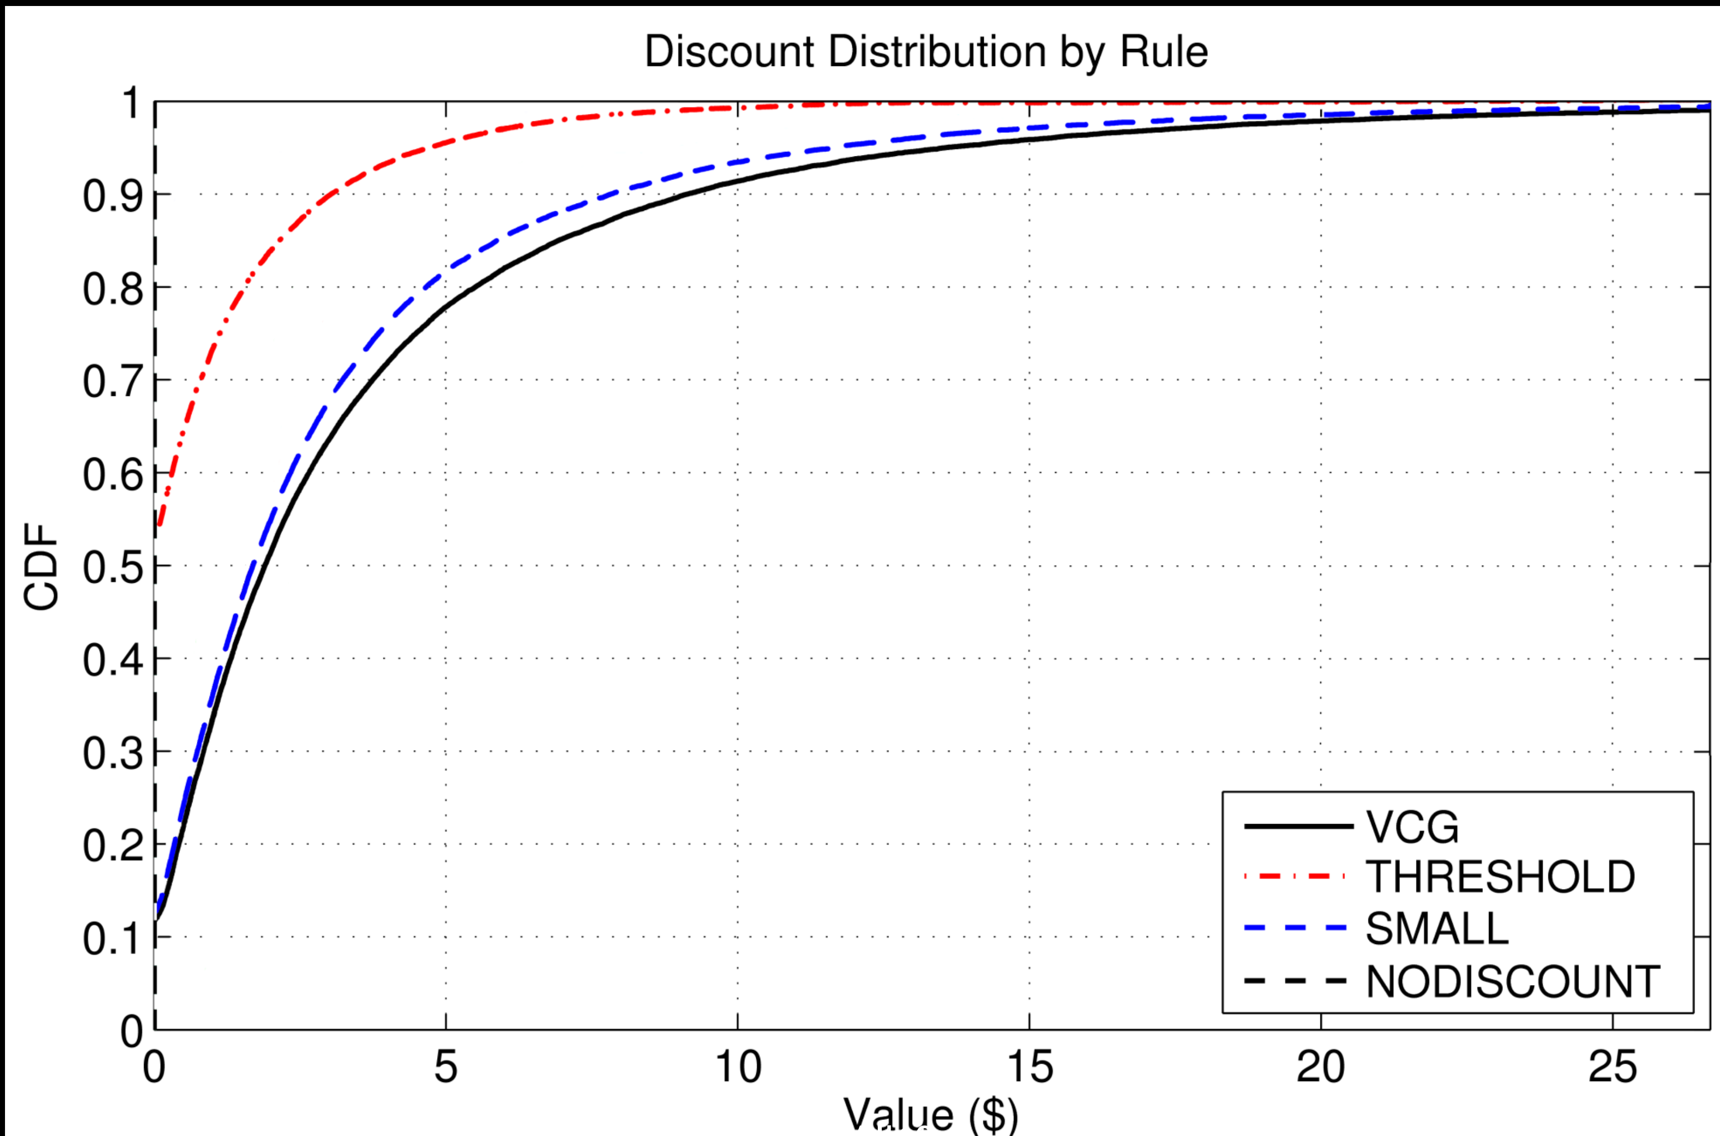

$$KLnorm(m) = \int_0^{\infty} \widehat{H}^*(\pi) log\left(\frac{\widehat{H}^*(\pi)}{\widehat{H}^m(\pi)}\right) d\pi$$

$$L_1(m) = \int_v ||\pi_+^*(v), \pi_+^m(v)||_1 \, g(v)dv \qquad (2)$$

$$L_1 \, norm(m) = \int_v ||\frac{\pi_+^*(v)}{V^*(v)}, \frac{\pi_+^m(v)}{V^*(v)}||_1 \, g(v)dv \qquad (3)$$

$$L_2(m) = \int_v ||\pi_+^*(v), \pi_+^m(v)||_2 \, g(v)dv \qquad (4)$$

Discriminative power of 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)

Other Approx SP Concepts

- SPITL: SP in a large market (Budish'09)
  - get best outcome in choice set
  - choice set becomes agent-independent in limit of continuum market
- Counting manipulations (Pathak & Sonmez'09)
  - (roughly) B manipulable by less agents in less instances than A
- Marginal gain (Erdil & Klemperer'09)
  - minimize $|\partial\pi_i (v_i, b_{-i}) / \partial v_i|$

Conclusion

- Incentive-efficient CEs are only know for simple settings (e.g., known single-minded bidders; one buyer, one seller, single unit)
- ICE = bidding language, winner determination, price feedback, proxy agents
- Payment design
  - "small" >> "threshold"
  - maximize # agents with zero regret
  - or, minimize divergence to VCG payoffs