Preference Elicitation and Preference Learning in Social Choice: New Foundations for Group Recommendation

Craig Boutilier
University of Toronto

(Joint work with Tyler Lu)

- Social choice: study of collective decision making
- *Aggregation of individual preferences* determines a *consensus outcome* for some population
  - Political representatives, committees, public projects,…
  - Studied for millennia, formally for centuries
- Increasing importance for low stakes domains…

Social Choice

Why *Computational* Social Choice

- Computational models/tradeoffs inherently interesting
  - Winner determination, manipulation/control, approximations, computational/communication complexity
- Decision making in multiagent systems
- Preference and rank learning in machine learning
  - Ready availability of partial preference data (web search data, ratings data in recommender systems, …)
- Complexity: combinatorial nature of alternatives

- Move to lower stakes, complex domains makes new demands on social choice
  - New models and decision criteria reflecting new uses
- Focus today: minimizing amount of information needed to come to *good* consensus choice
  - Robust decision making with partial rankings/votes
  - Incremental elicitation of voter preferences
  - Exploiting distributional information
  - *Learning probabilistic models of population preferences*
  - Extensions:
    - *Combinatorial alternative spaces*
    - *Voting on social networks*

Our Agenda

Overview**

- Social Choice: Main Concepts
- Regret-based Vote Elicitation
  - Minimax regret (robustness criterion) for partial vote
  - Polytime computation of MMR for certain rules
  - Elicitation of voter preferences using MMR
- Optimal One-round Elicitation Protocols
- Brief: Learning Mallows Models of Population Prefs
- *Next Steps*

Social Choice: Basic Framework

- *Alternative set* $A = \{a_1, \ldots, a_m\}$ 
- *Voters* $N = \{1..n\}$, each with preferences over A
- *Vote* v_i of voter i: a linear ordering (permutation) of A
- *Profile* is collection of votes $\mathbf{v} = (v_1, \ldots, v_n)$
- *Winner:* alternative maximizing "consensus"
  - or sometimes a *consensus ranking*

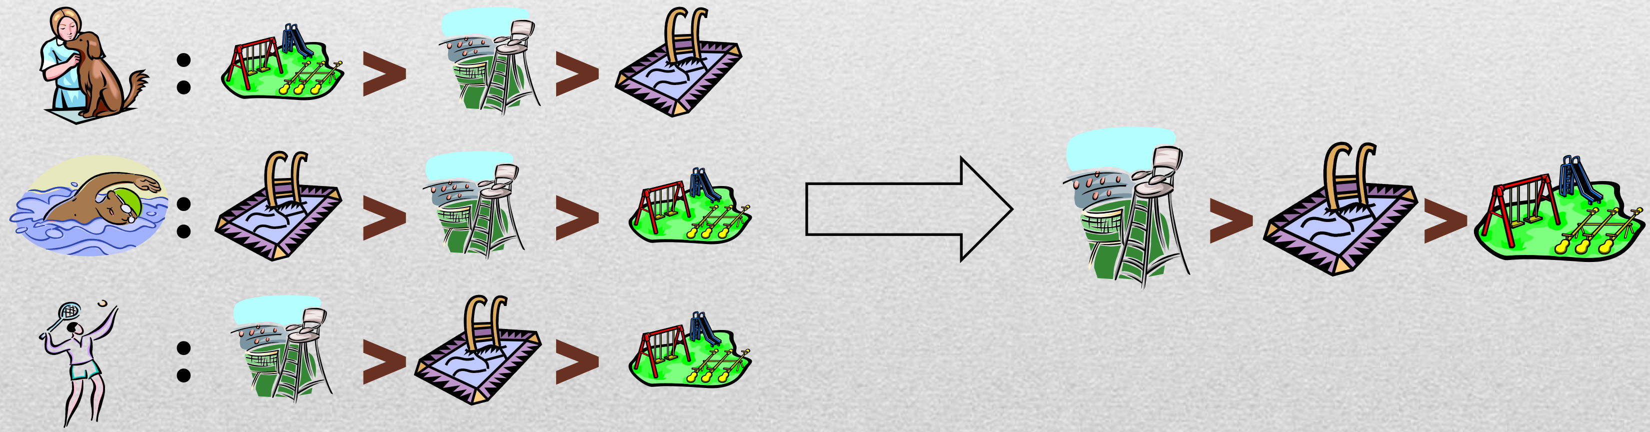

Voting Rules

- *Voting rule $r: V \rightarrow A$* selects a winner given a profile
- *Plurality:* winner *a* with most 1st-place votes
  - voters needn't provide full ranking
- *Positional scoring:* Assign *score α* to each rank position with

  $$\alpha(1) \geq \alpha(2) \geq \ldots \alpha(m)$$

  - *Borda count* well-known: $\alpha = <m-1, m-2, \ldots, 0>$
  - Winner: *a* with max sum of scores: $\sum_i \alpha(v_i(a))$
  - Plurality, k-approval, k-veto special cases

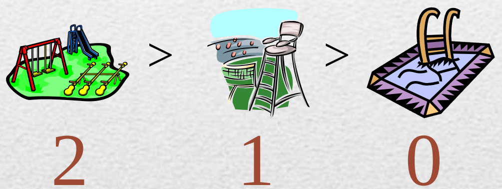

- *Maxmin Fairness (egalitarian):*
  - Score of *a* is $min\ \{i: m - v_i(a)\}$
  - Choose *a* with highest score

- Many other rules: Copeland, maximin, Bucklin, etc.

- Most voting rules have "natural" scoring functions $s(a, \mathbf{v})$
- $s(a, \mathbf{v})$ measures "quality" of alternative a given profile $\mathbf{v}$
- Rule r chooses $r(\mathbf{v}) \in argmax \ \{s(a, \mathbf{v}) : a \in A\}$

Score-based Voting Rules

Vote Elicitation [Lu, B. IJCAI-11]

- Use of complex (rank-based) voting schemes rare
  - Cognitive complexity, communication costs, *monetary costs*
- Elicitation of *partial votes* could ease this burden
  - Find *relevant* comparisons… or even *approximate* winners

- **Voting Protocol with Approximation:** Ask a few queries of voters: if close enough, stop; otherwise ask a few more; continue until satisfied

- Theoretically, relevance won't save much:
  - Communication complexity *O(nm log m)* for Borda, etc. **[CS EC-05]**
  - This doesn't mean practical savings are not possible!

Partial Vote Profiles

- *Partial vote* p_i of voter *i*: consistent set of pairwise comparisons of form $a_j \succ a_k$
  - Captures most natural constraints: paired comp, top-k, etc.
- *Partial profile* $\mathbf{p} = (p_1, \ldots, p_n)$
- *Completions* $C(p_i)$, $C(\mathbf{p})$: set of votes extending p_i, $\mathbf{p}$

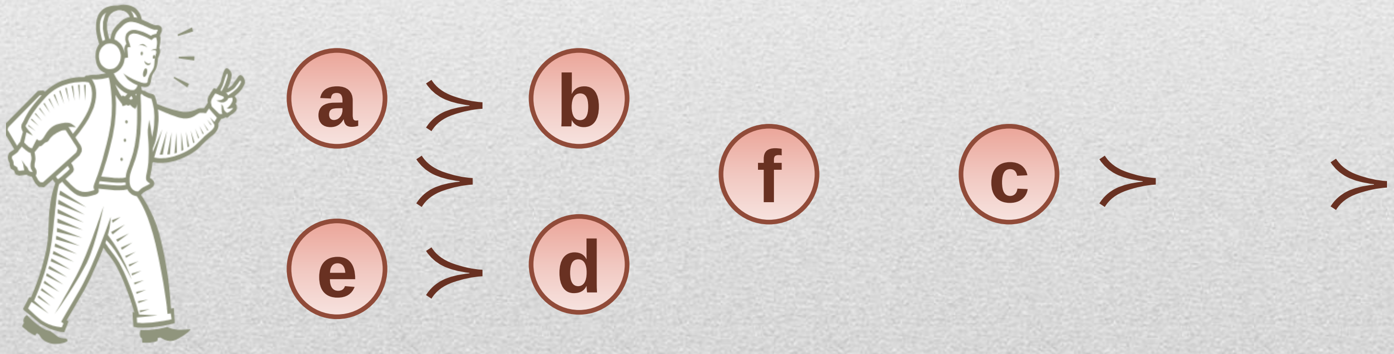

Robust Winner Determination

- In general, may want to decide given a partial profile
  - Robustness criteria rarely discussed in social choice
- We propose *minimax regret* to determine winners

$$Regret(a,\mathbf{v}) = max_{a' \in A} s(a',\mathbf{v}) - s(a,\mathbf{v})$$
$$= s(r(\mathbf{v}),\mathbf{v}) - s(a,\mathbf{v})$$

$$PMR(a,a',\mathbf{p}) = max_{\mathbf{v} \in C(\mathbf{p})} s(a',\mathbf{v}) - s(a,\mathbf{v})$$
$$MR(a,\mathbf{p}) = max_{\mathbf{v} \in C(\mathbf{p})} Regret(a,\mathbf{v})$$
$$= max_{a' \in A} PMR(a,a',\mathbf{p})$$

Adversarial choice

$$MMR(\mathbf{p}) = min_{a \in A} MR(a,\mathbf{p})$$
$$a^*_{\mathbf{p}} \in \operatorname*{argmin}_{a \in A} MR(a,\mathbf{p})$$

Best response

Minimax Regret: Illustration

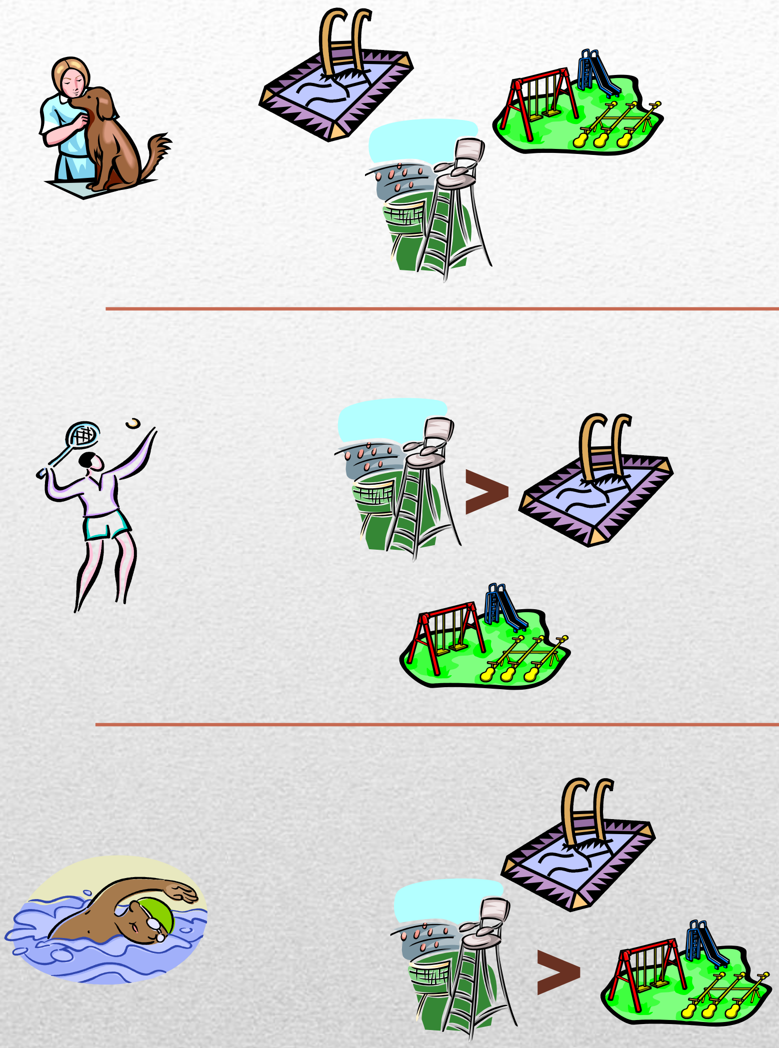

Proposed Winner: Tennis

Minimax Regret: Illustration (Borda)

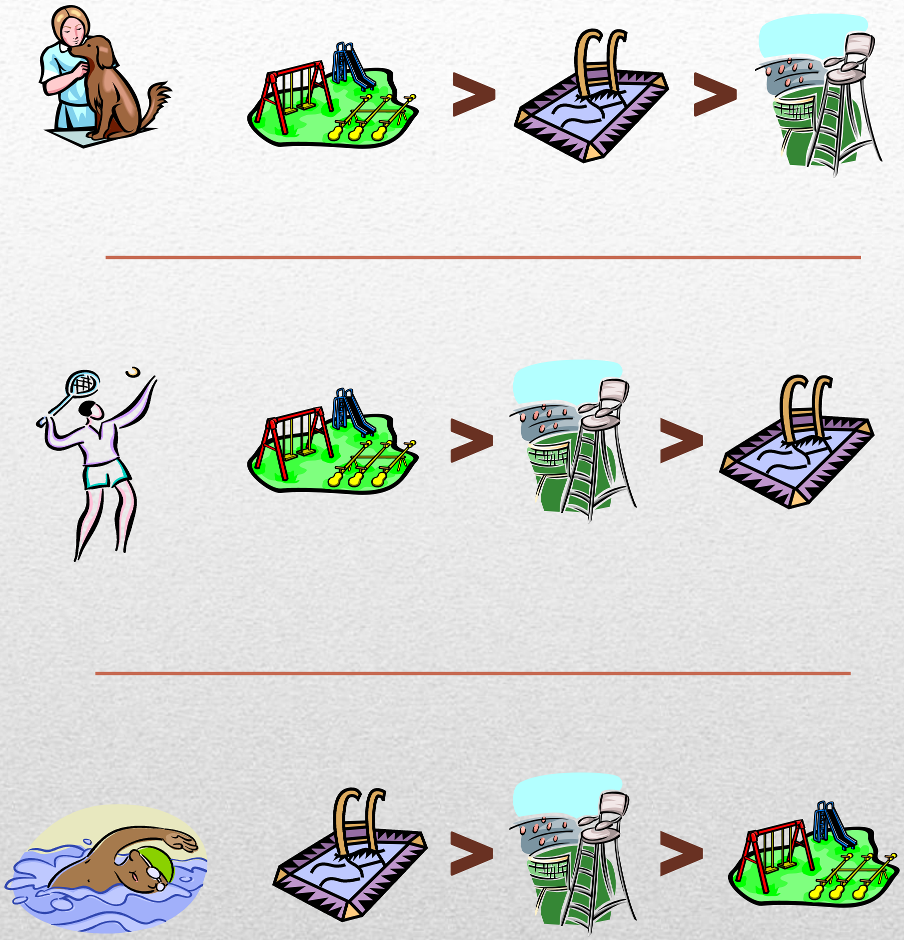

Proposed Winner: Tennis

Borda Score(Tennis) = 2
Borda Score (Park) = 4
Max Regret(Tennis) = 2 *(4-2)*

Proposed Winner: Pool

Minimax Regret: Illustration (Borda)

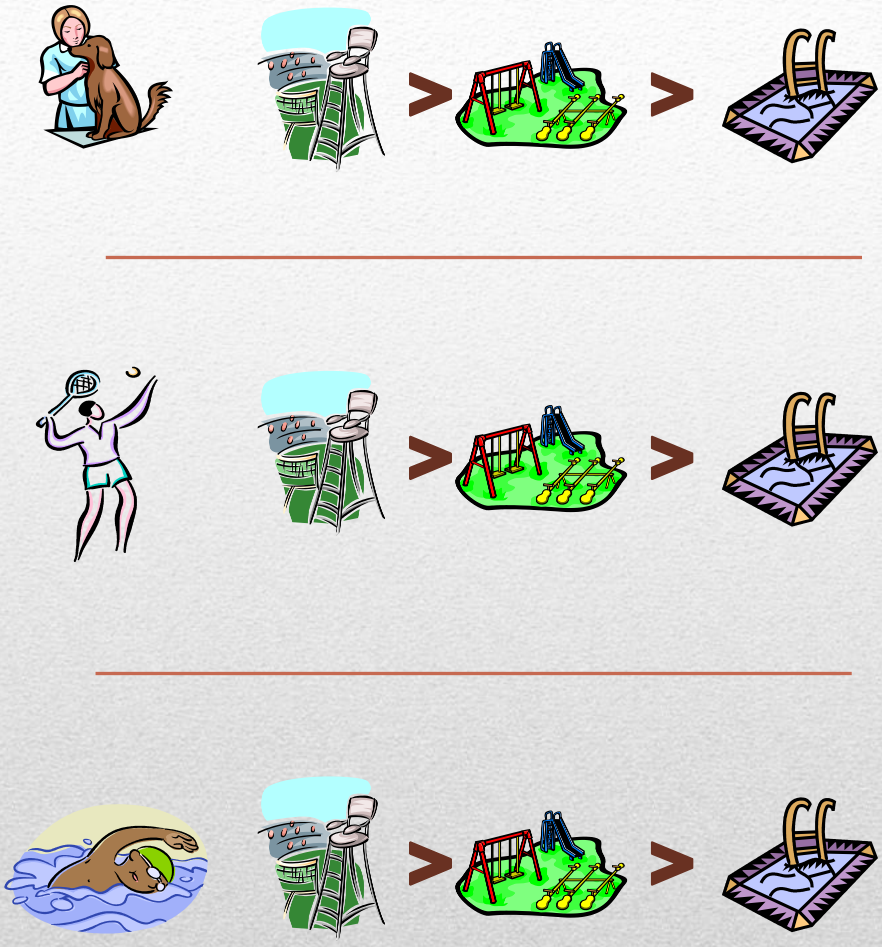

Proposed Winner: Tennis

Borda Score(Tennis) = 2
Borda Score (Park) = 4
Max Regret(Tennis) = 2 *(4-2)*

Proposed Winner: Pool

Borda Score(Tennis) = 6
Borda Score (Pool) = 0
Max Regret(Pool) =6 *(6-0)*

Minimax Optimal: Tennis
Minimax Regret: 2

- MMR offers a natural robustness criterion
  - candidate with *tightest* error bounds (loss wrt optimal)
  - provably optimal if MMR=0
- Contrast with maximin
  - provides quality guarantee, not optimality guarantee
- Contrast with Bayesian methods, which have/are:
  - need for a prior
  - no (worst-case) guarantees
  - computationally difficult (even to approximate)

Why Minimax Regret

- *MMR(**p**)=0* iff winner $a_{\mathbf{p}}^*$ is a necessary co-winner
- **Obs:** MMR computation at least as hard as NecCo-Win
- **Obs:** MMR-winner may not be a *possible winner*
  - In fact, *all* possible winners may have high max regret

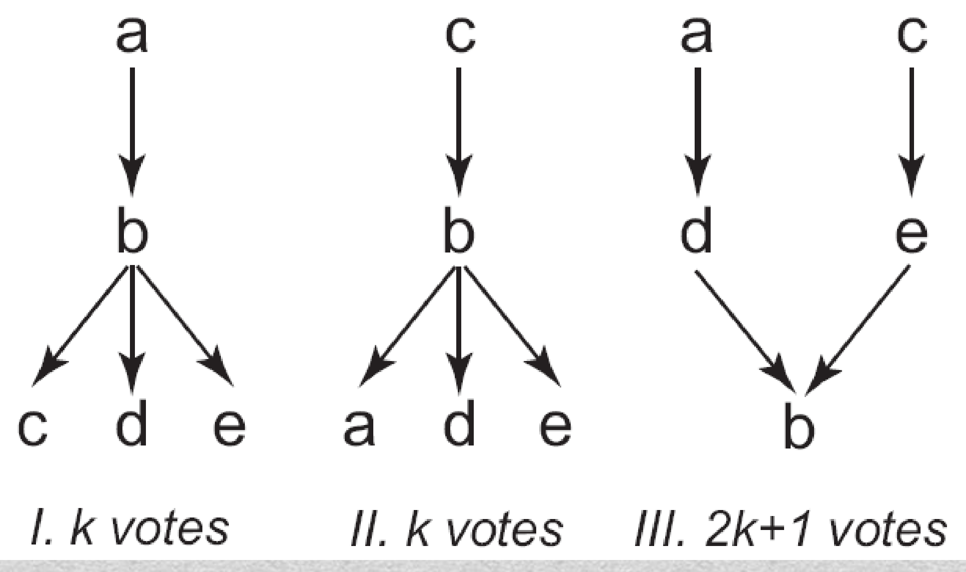

Assume 2-approval:
- Only *a, c* are PWs: one has score at least *2k+1*, while *b* has score *2k*
- *MR(b) = k+1*
- *MR(a) = MR(c) = 2k+1*

MR of *a, c* twice that of *b*

Properties of Minimax Regret Solution**

Computing Minimax Regret

- MMR for many problems often specified as an IP
  - Problematic for voting: too many voters/variables
- Instead, compute *PMR (a, w, **p**)* for all m^2 pairs *(a,w)*
  - Then $MMR(\mathbf{p}) = \min_a \max_w PMR(a, w, \mathbf{p})$

PMR	a	b	c	MR
a	0	2	2	2
b	2	0	6	6
c	5	3	0	5

- PMR can be computed in polytime for many rules
  - find *worst case completion* of each voter's partial vote p_i; can usually be done independently for each voter
    - Xia, Conitzer (AAAI08) use similar ideas for necessary winners
  - we illustrate with the Borda rule

Computing Minimax Regret**

- We illustrate with Borda (positional) scoring
- Positional: additively decomposable: $s(a, \mathbf{v}) = \sum_i s(a, v_i)$
- Thus PMR decomposable: complete each p_i independently

$$\begin{aligned} Regret(a, w, \mathbf{v}) &= s(w, \mathbf{v}) - s(a, \mathbf{v}) \\ &= \sum_i s(w, v_i) - \sum_i s(a, v_i) \\ &= \sum_i [s(w, v_i) - s(a, v_i)]. \end{aligned}$$

$$\begin{aligned} PMR(a, w, \mathbf{p}) &= \max_{\mathbf{v} \in C(\mathbf{p})} s(w, \mathbf{v}) - s(a, \mathbf{v}) \\ &= \sum_i \max_{v_i \in C(p_i)} s(w, v_i) - s(a, v_i). \end{aligned}$$

- Fix partial vote *p*: proposed alternative *a* and adversarial witness *w* stand in only one of three relations in *p*

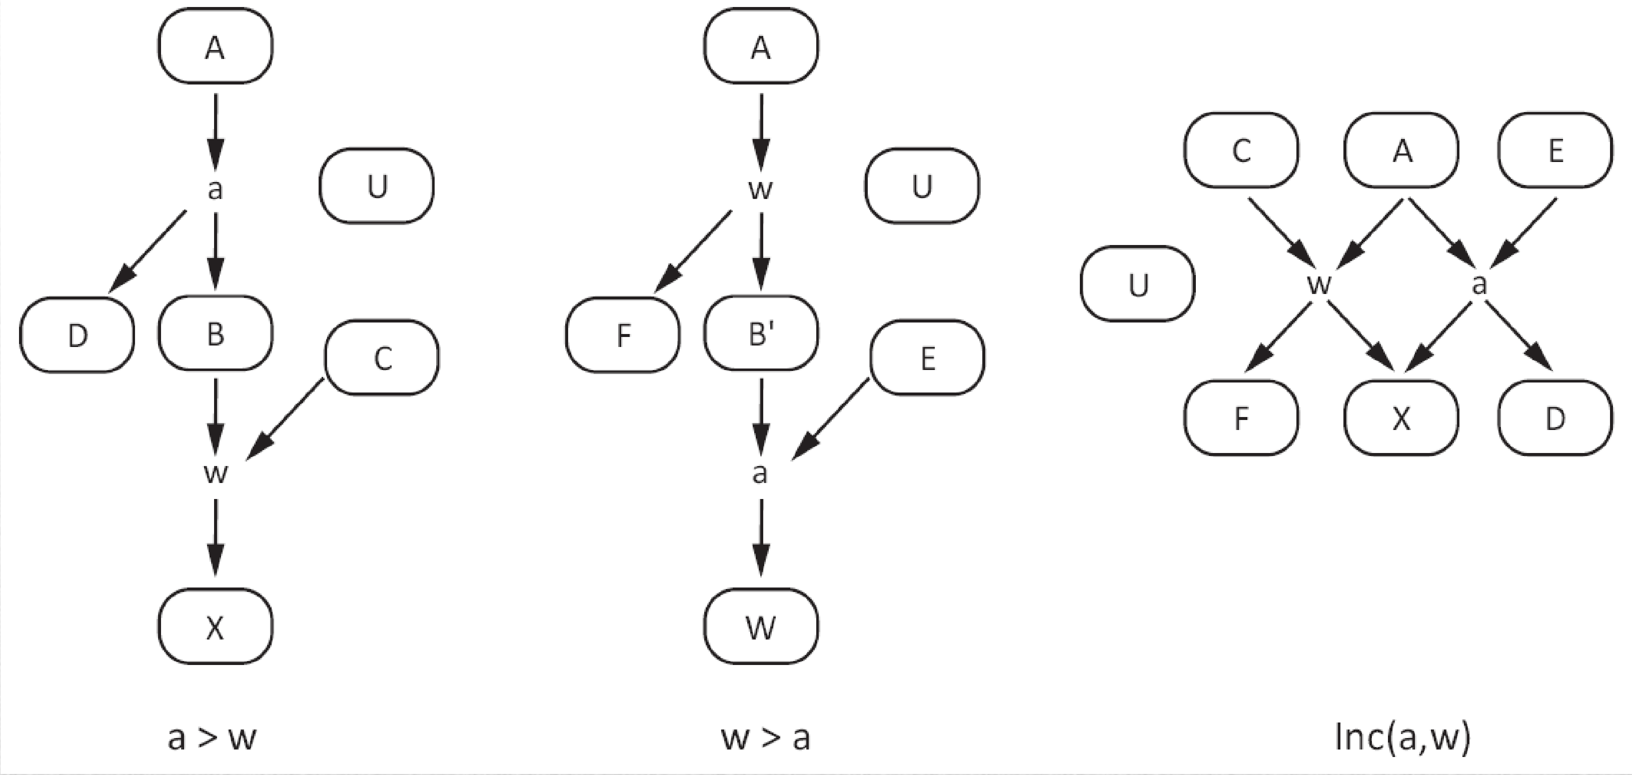

Computing Minimax Regret

Computing Minimax Regret

- Case 2: Maximize *PMR(a,w)* in only "one" way:

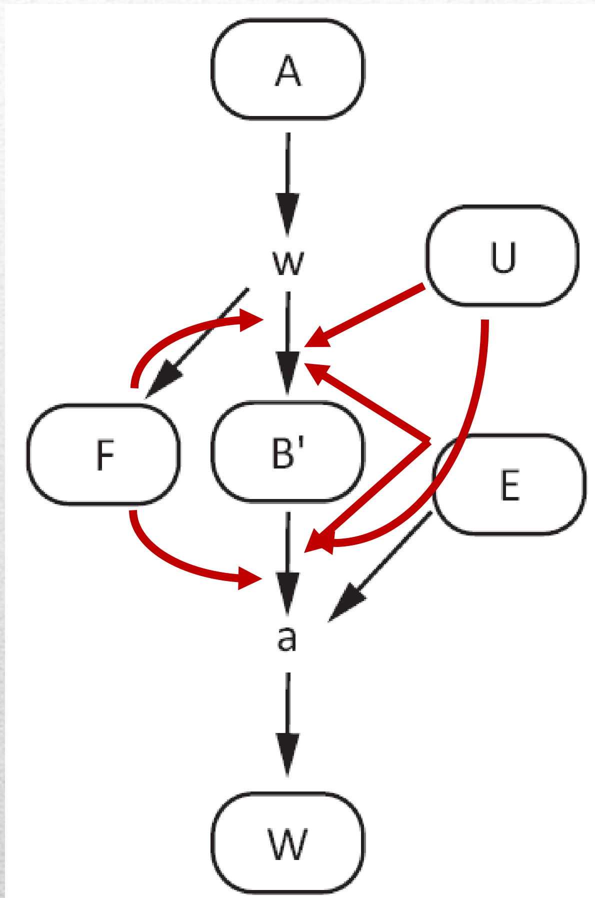

$$PMR(a,w)$$
$$= |B' \cup F \cup E \cup U| + 1$$

$$= m - (|A \cup W| + 1)$$

- Case 1: Maximize *PMR(a,w)* in only "one" way:

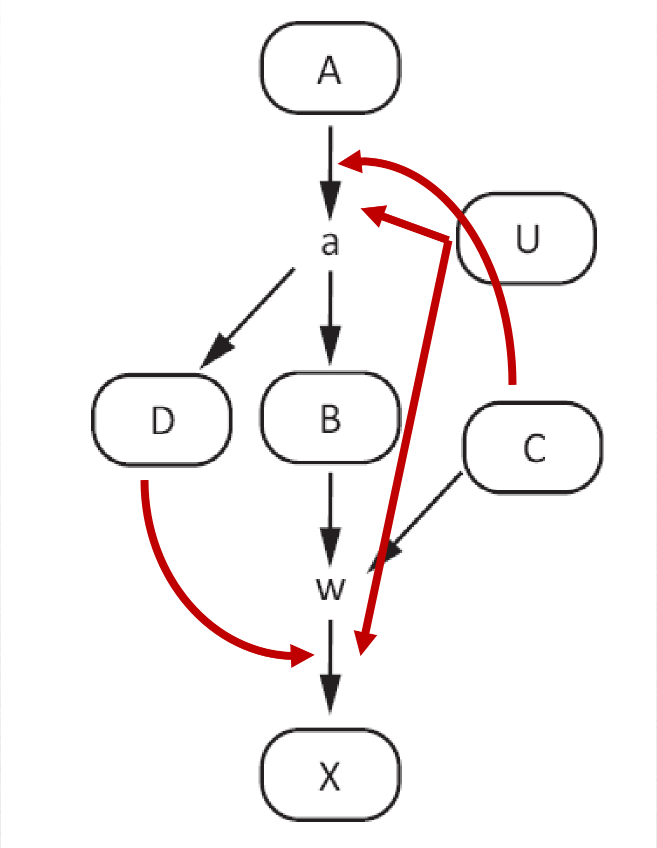

PMR(a,w) = −(|B| + 1)

Computing Minimax Regret**

Computing Minimax Regret**

- Case 3: Maximize *PMR(a,w)* in only “one” way:

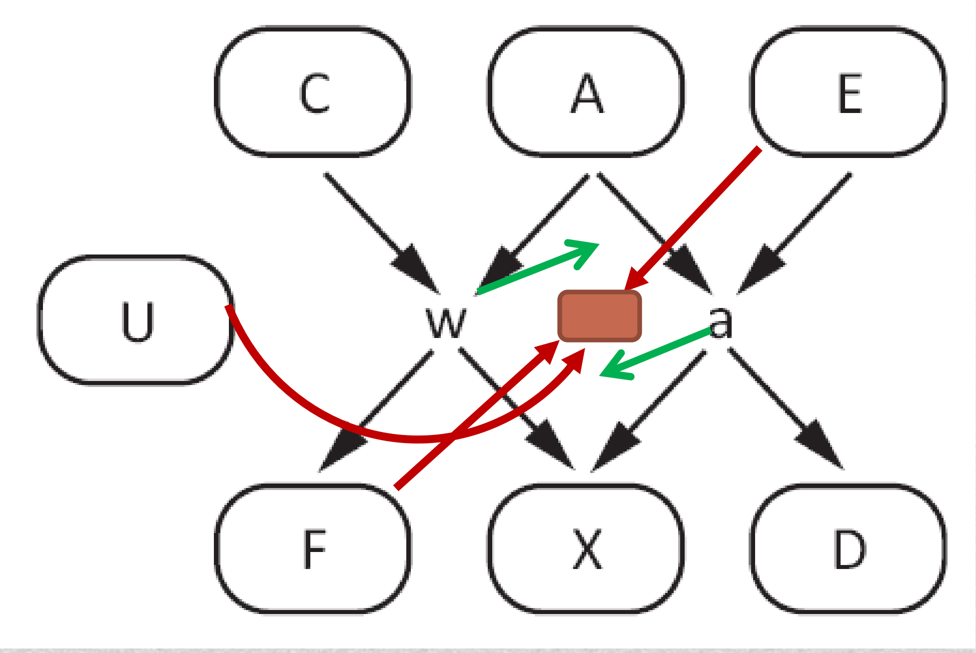

PMR(a,w) =
|F ∪ E ∪ U| +1

- Similar analysis: other positional scoring rules
- Similar approach for non-decomposable scoring rules
- Max regret computation is polytime for:
  - Positional scoring rules
  - Egalitarian (maxmin fairness)
  - Bucklin
  - Maximin

Computing Minimax Regret

Regret-based Vote Elicitation

- If *MMR(**p**)* too high, refine knowledge of voter preferences
- Current Solution Strategy (CSS):
  - Use MMR solution *(a*,w)* to generate query: if we don't reduce *PMR(a*,w),* MMR will not be reduced
  - So find some voter *i* with vote p_i and ask query with potential to reduce advantage of *w* over *a** in $C(p_i)$
  - For each voter, queries considered depend on structural properties of partial vote (whether Case 1, 2, 3; and size of sets)

Regret-based Vote Elicitation

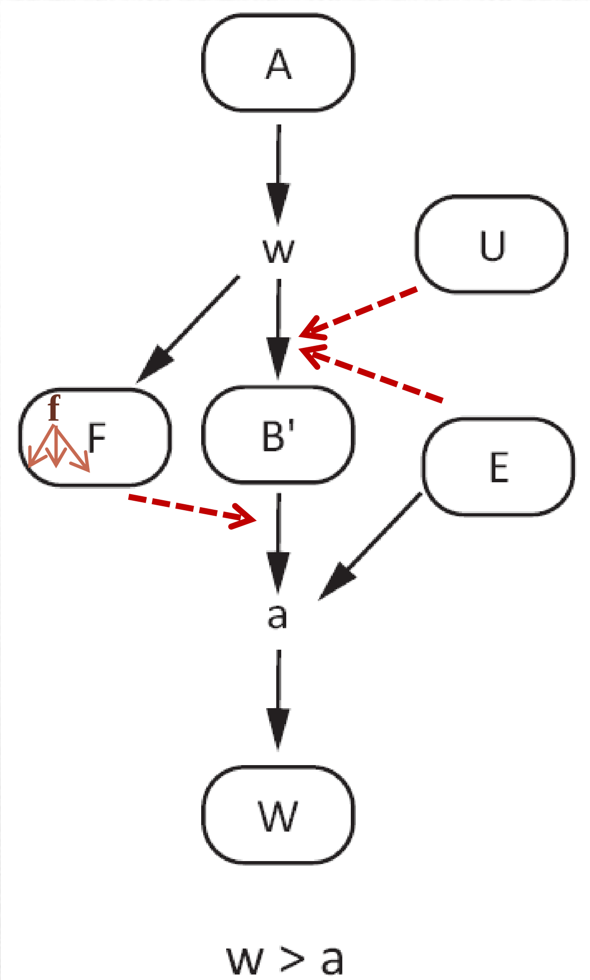

Case 2: four reasonable query types

- $a \succ f$ for some $f \in F$
  - Max potential: f at "top" of large group
- $a \succ u$ for some $u \in U$
  - Max potential: u at "top" of large group
- $e \succ w$ for some $e \in E$
  - Max potential: e at "bottom" of large group
- $u \succ w$ for some $u \in U$
  - Max potential: u at "bottom" of large group

Note: if MMR>0, one of U,E,F nonempty for some voter (or sets in cases 1, 3)

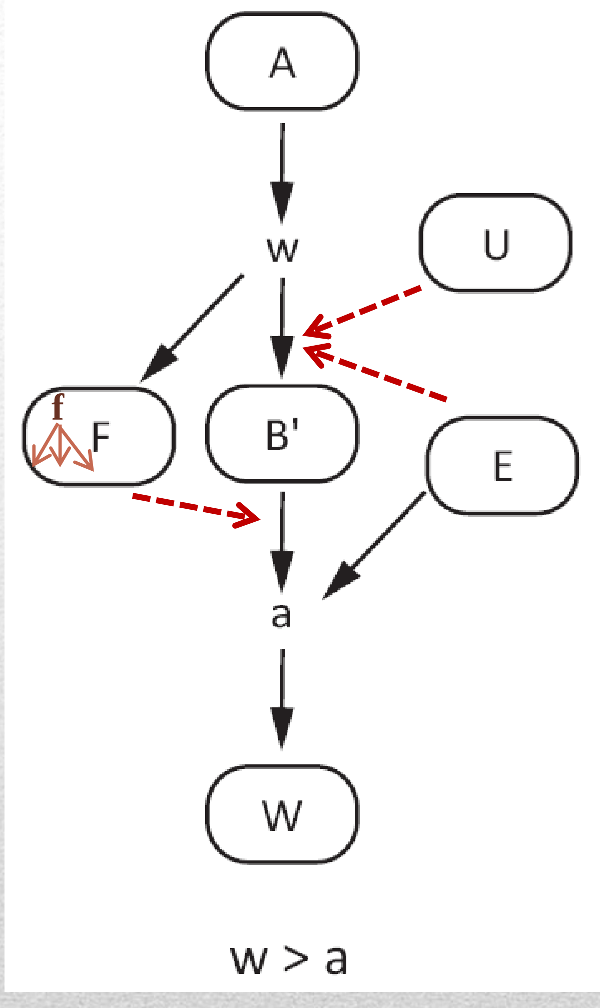

Case 2: four reasonable query types

- $a \succ f$ for some $f \in F$
  - Max potential: f at "top" of large group
- $a \succ u$ for some $u \in U$
  - Max potential: u at "top" of large group
- $e \succ w$ for some $e \in E$
  - Max potential: e at "bottom" of large group
- $u \succ w$ for some $u \in U$
  - Max potential: u at "bottom" of large group

- *Note: if MMR>0, one of U,E,F nonempty for some voter (or sets in cases 1, 3)*

Regret-based Vote Elicitation

Vote Elicitation: Experiments*

- Intuitions behind pairwise CSS can be generalized to *top-t queries* (only pick voter, not alternative pair)

- Compare CSS to two strategies
  - Volumetric: choose voter/candidate-pair which introduces greatest number of new paired comparisons

$$Vol(p_k) = \max_{a_i,a_j} \min\{|tc(p_k \cup \{a_i \succ a_j\})|, tc(v \cup \{a_j \succ a_i\})\}$$

  - Rand: random voter/candidate pair

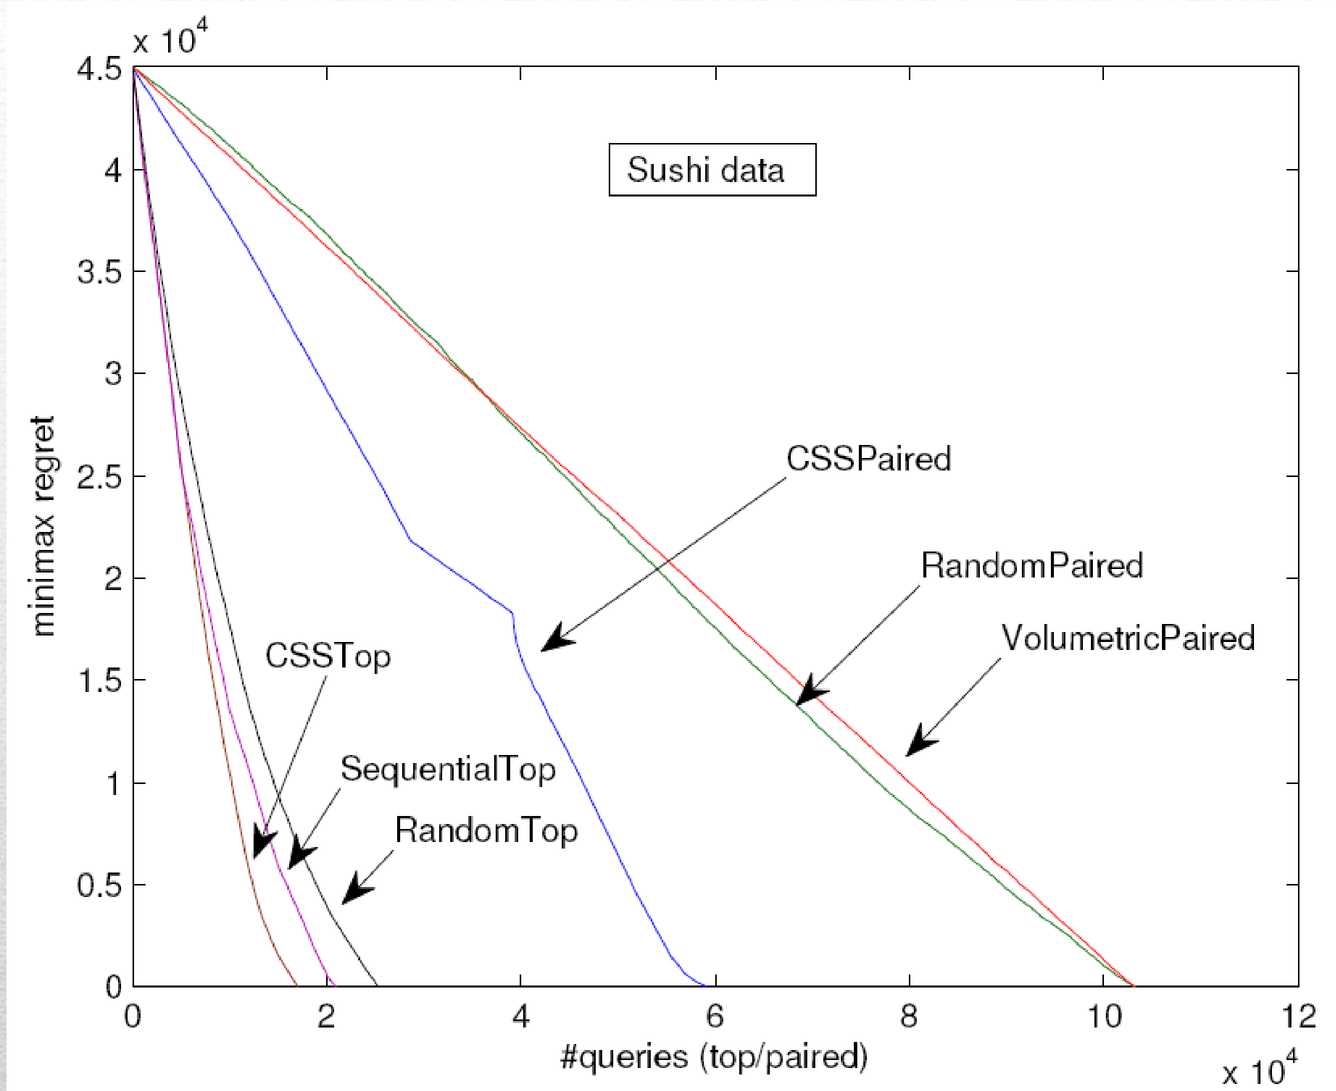

Sushi: 5000 rankings of 10 varieties of sushi

Vote Elicitation: Sushi

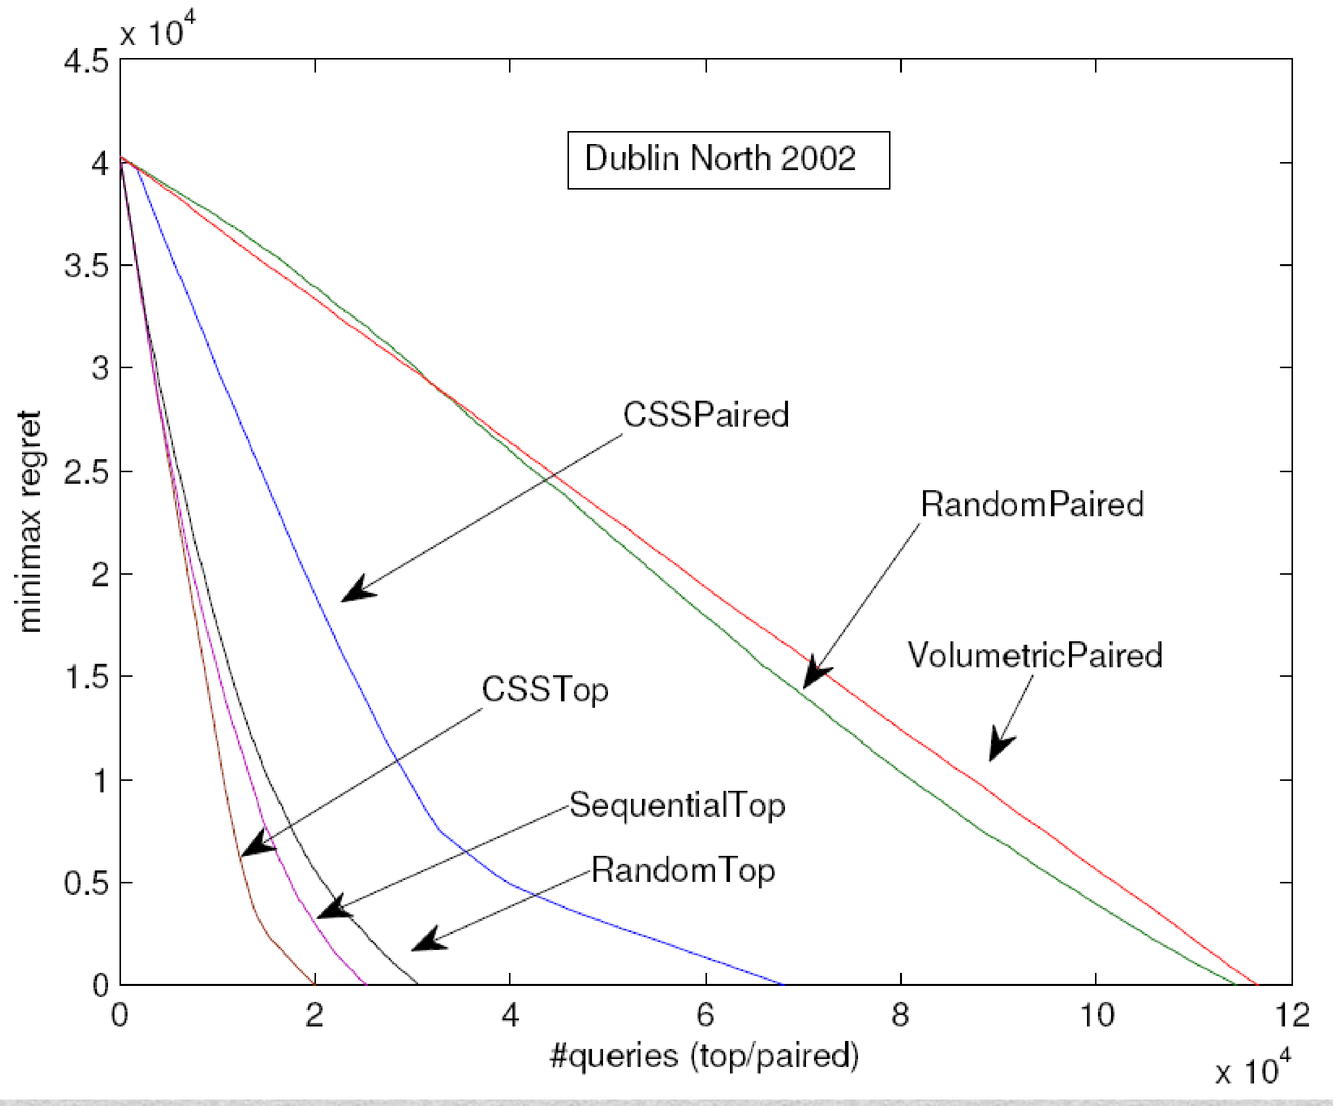

Irish: 2002 electoral data (Dublin North); 3662 rankings over 12 candidates

Vote Elicitation: Dublin North 2002

- Let $d(r,\sigma)$ denote Kendall-tau distance
  - Number of pairwise inversions (swaps) between r, σ
- Let σ be some central/modal ranking
- *Mallows ϕ-model* (with dispersion ϕ)specifies $P(r)$:

$$P(r) = P(r \mid \sigma, \phi) = \frac{1}{Z}\phi^{d(r,\sigma)}$$

- If $\phi = 1$, P is uniform (IC); as $\phi \rightarrow 0$, P concentrates on σ
- Unimodal nature of model inflexible; but *mixtures of Mallows models* can reasonably capture certain types of population preferences

Mallows Models

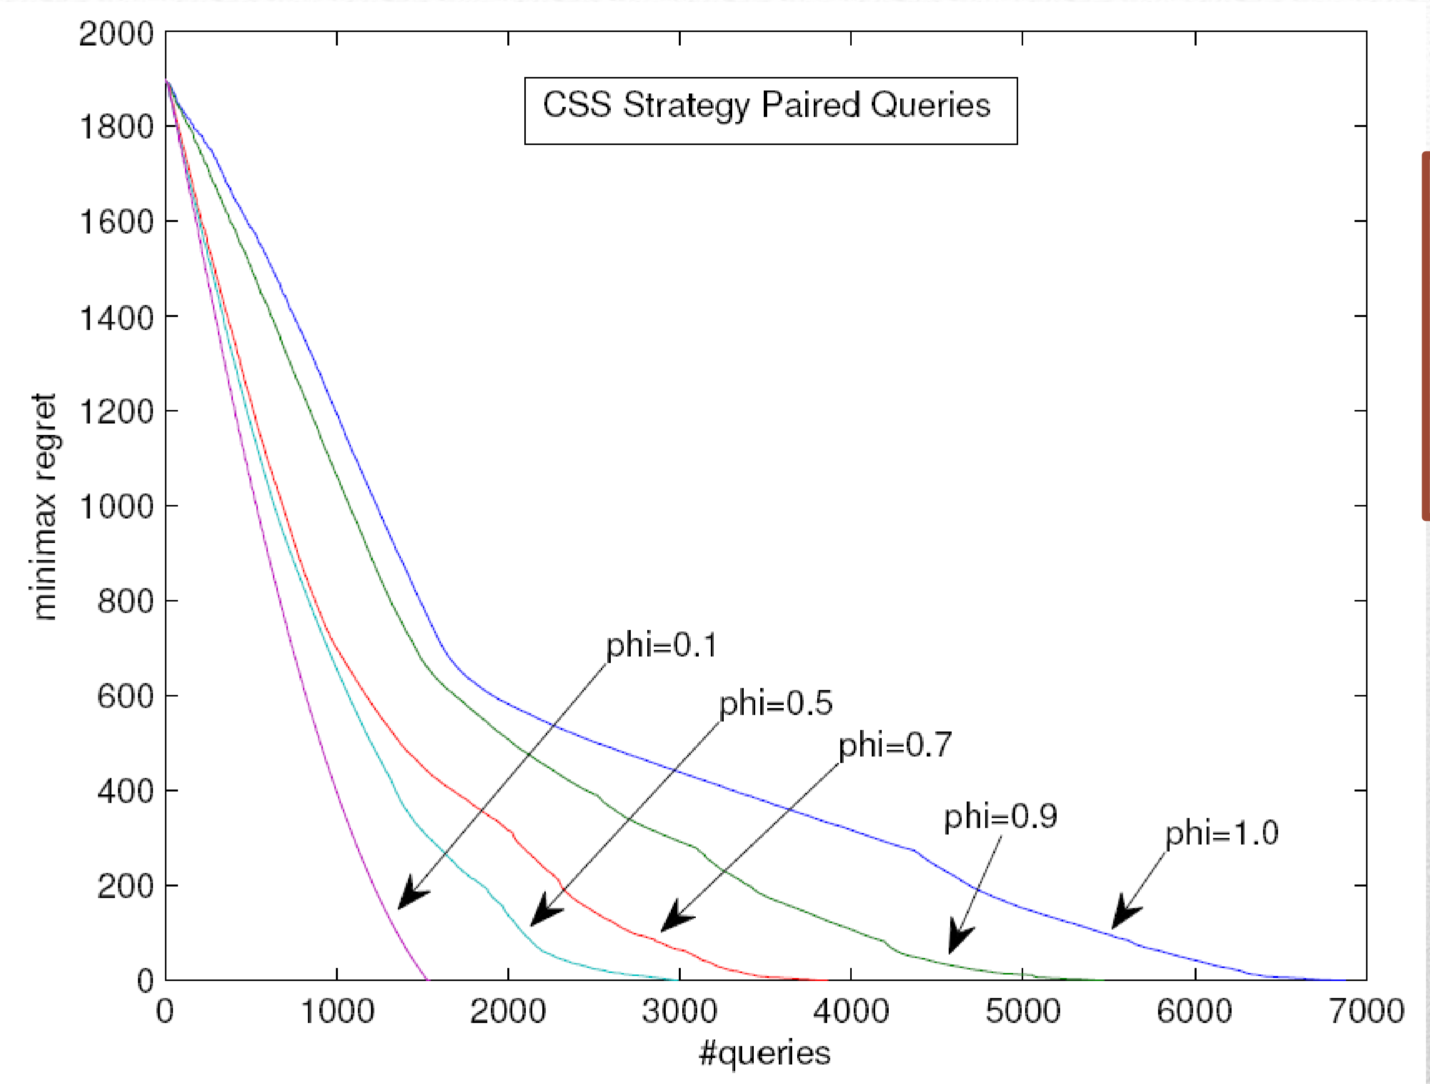

Mallows: 100 random rankings over 20 items; vary dispersion ϕ

Vote Elicitation: Mallows

Summary of Results

- MMR=0 after *k* paired comparisons per voter
  - *Sushi:* CSS 11.82; Vol 20.64; Rand 20.63; MergeSort 25
  - *Irish:* CSS 18.57; Vol 31.82; Rand 31.22; MergeSort 33

- MMR=0 after *k* top-t queries per voter
  - *Sushi:* CSS 3.40; Vol 4.18; Rand 5.50
  - *Irish:* CSS 5.47; Vol 6.91; Rand 8.38

- Anytime performance better for CSS as well
  - E.g., reach 18% of initial regret on Irish data set after only 5.82 queries (vs. 25.77 Vol; 24.03 Rand)

Single vs. Multi-round Elicitation

- Fully sequential elicitation often not practical
  - *Tradeoff:* quality, information elicited, rounds/interruption
    - see Kalech et al. **[JAAMAS 2011]**
- Reduce interruption cost by using coarser “rounds”
  - E.g., ask each voter for their *top k* candidates
  - Stop if MMR low enough
  - Otherwise select a few voters and ask for their *next k’* candidates; etc.
- *Suitable choice of k* balances the three criteria

- General framework for addressing tradeoffs
- Focus on optimizing *single-round protocols*
  - for one round of elicitation, what is trade off between information elicited (k) and minimax regret?

- Requires a probabilistic model *Pr* of voter preferences
  - weak guarantees otherwise (hard to *predict* MMR)
- Our goal: find minimal k s.t. $Pr(MMR < \varepsilon) > 1 - \delta$
  - *regret tolerance* ε
  - *confidence* δ

Optimizing Single-round Protocols [Lu, B. ADT-11]

Exploiting Distribution: Sampling

- Many models of ranking distributions:
  - *Mallows, Plackett-Luce, Bradley-Terry, impartial culture, …*
  - in principal, can derive analytical results for each
- We propose an empirical (sampling) methodology
  - sample t vote profiles
    - learned model, generative process, subsample data sets
  - compute MMR for each profile and for each $k < m-1$
  - use empirical distribution over MMR to determine suitable k achieves desired $MMR < \varepsilon$ with desired probability $Pr > 1- \delta$

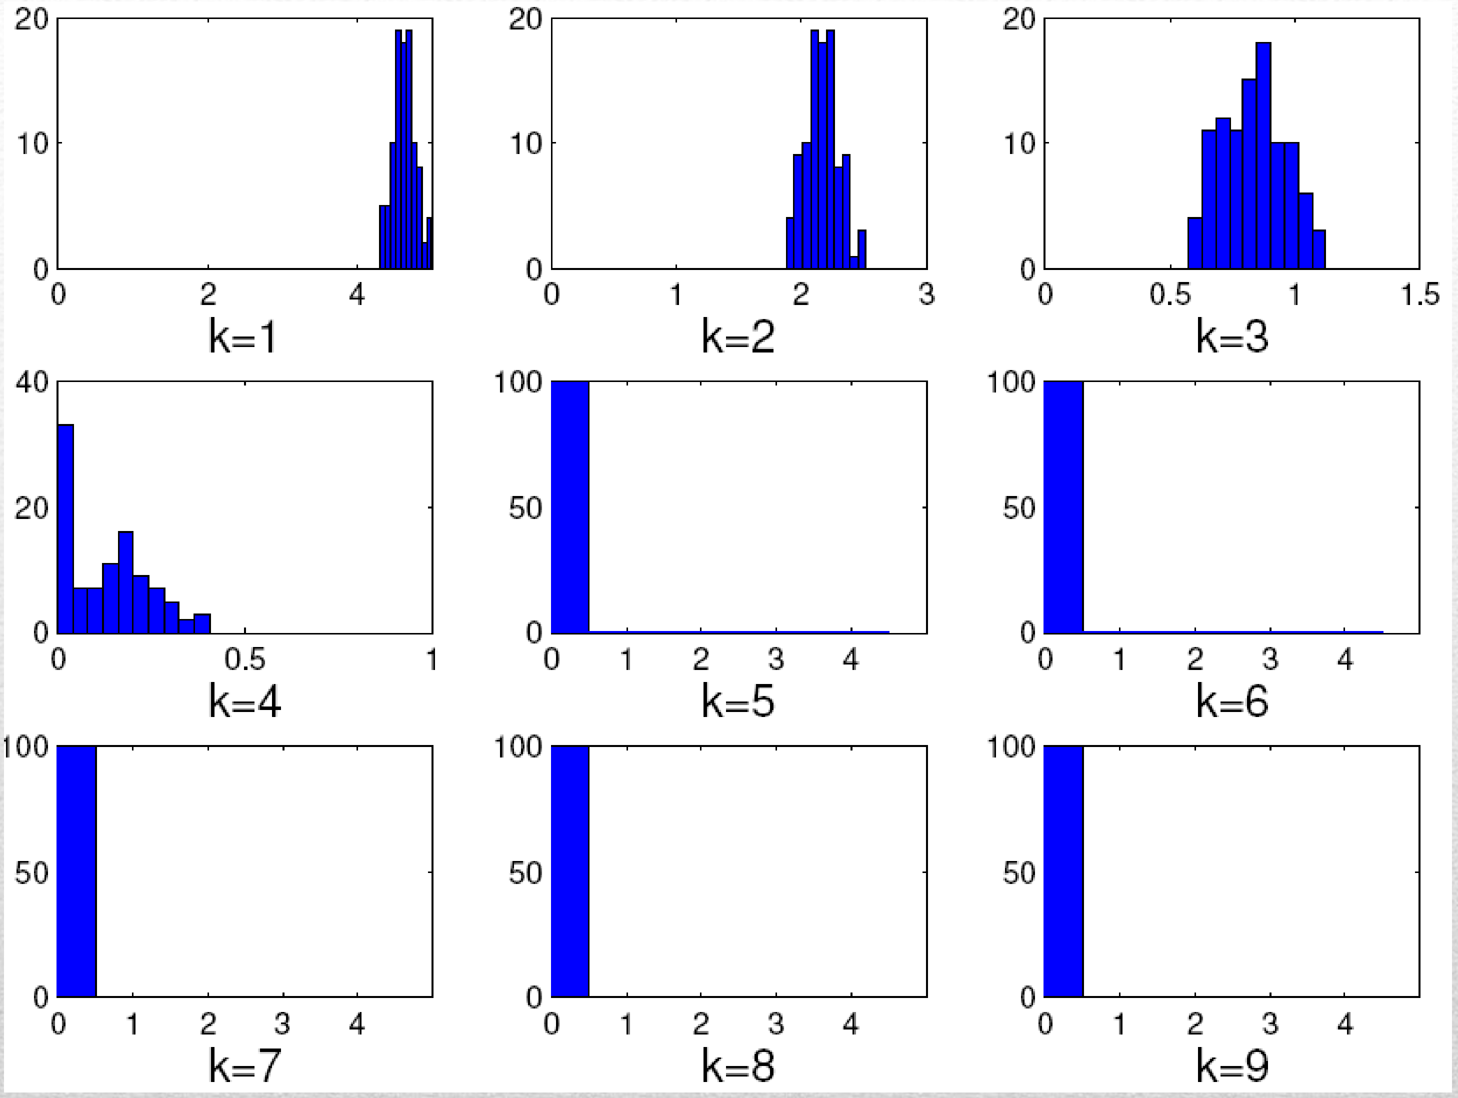

MMR Histograms: Mallows (m=10, n=1000, ϕ=0.6, Borda)

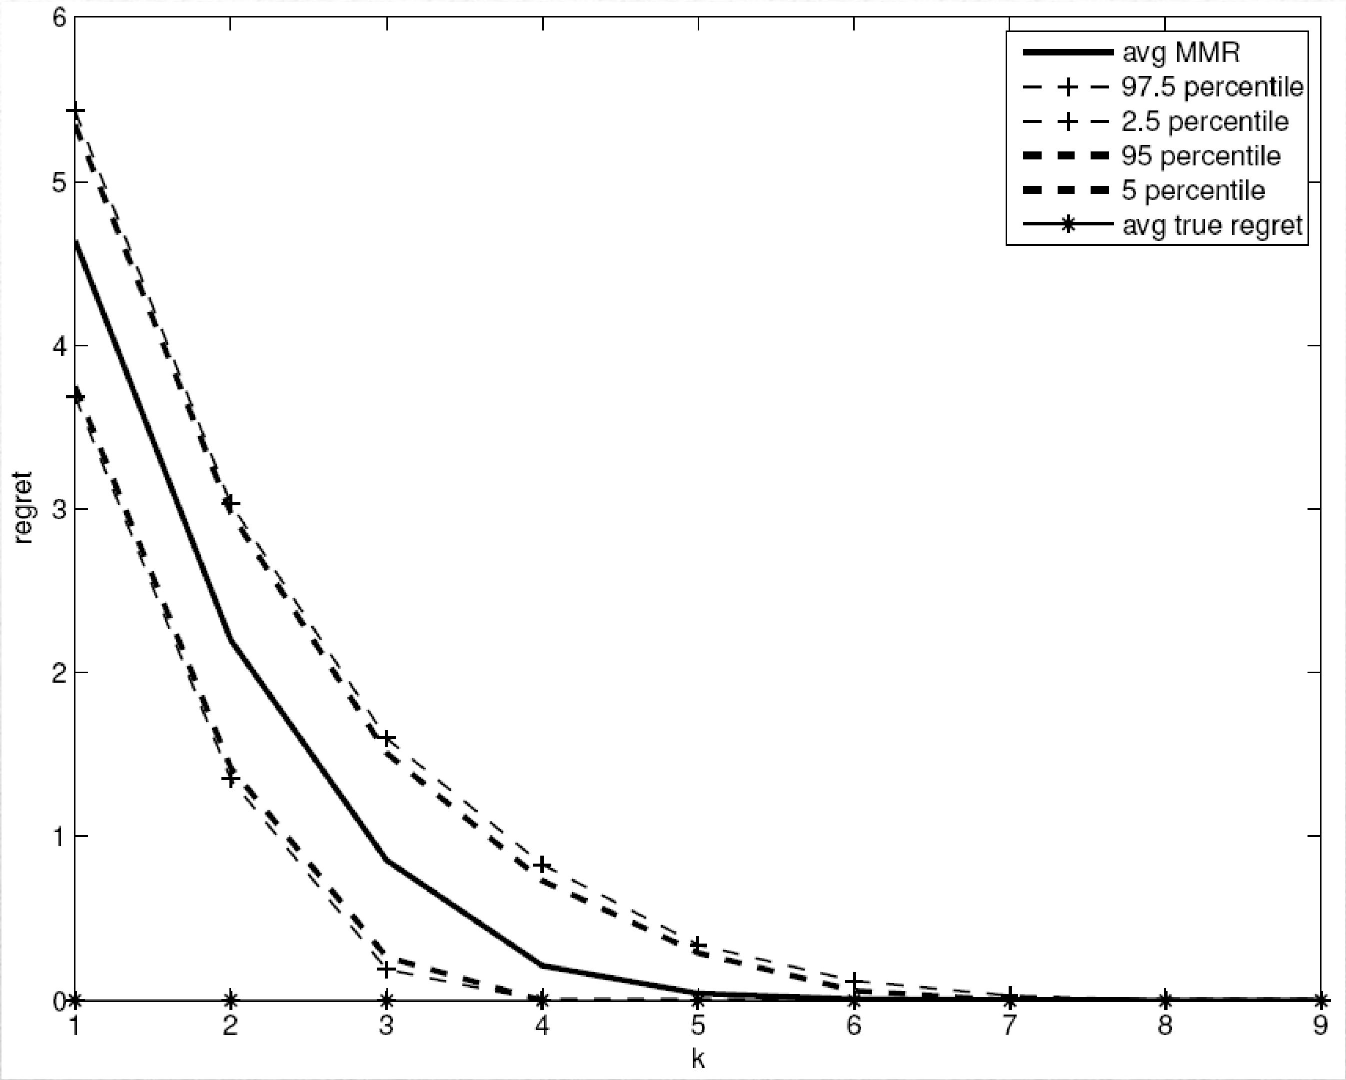

MMR Confidence Plot: Mallows (m=10, *n=100*, ϕ=0.6, Borda)

- One may use methodology purely heuristically
  - actual MMR (after elicitation) can suggest further queries
- Theoretical sample complexity bounds possible
  - assume sampling accuracy ξ and sampling confidence η
  - with t sampled profiles, where: $$t \geq \frac{1}{2\xi^2} \ln \frac{2(m-2)}{\eta}.$$
  - output min $\hat{k}$ satisfying: $$\hat{q}_k \equiv \frac{|\{i \leq t : MMR(\mathbf{p}_i[k]) \leq \varepsilon\}|}{t} > 1 - \delta - \xi$$

Theorem 1. *Let $\varepsilon, \delta, \eta, \xi > 0$. If sample size t satisfies Eq. 4, then for any preference profile distribution P, with probability $1 - \eta$ over i.i.d. samples $\mathbf{v}_1, \ldots, \mathbf{v}_t$, we have: (a) $\hat{k} \leq k^*$; and (b) $P[MMR(\mathbf{p}[\hat{k}]) \leq \varepsilon] > 1 - \delta - 2\xi$.*

Sample Complexity

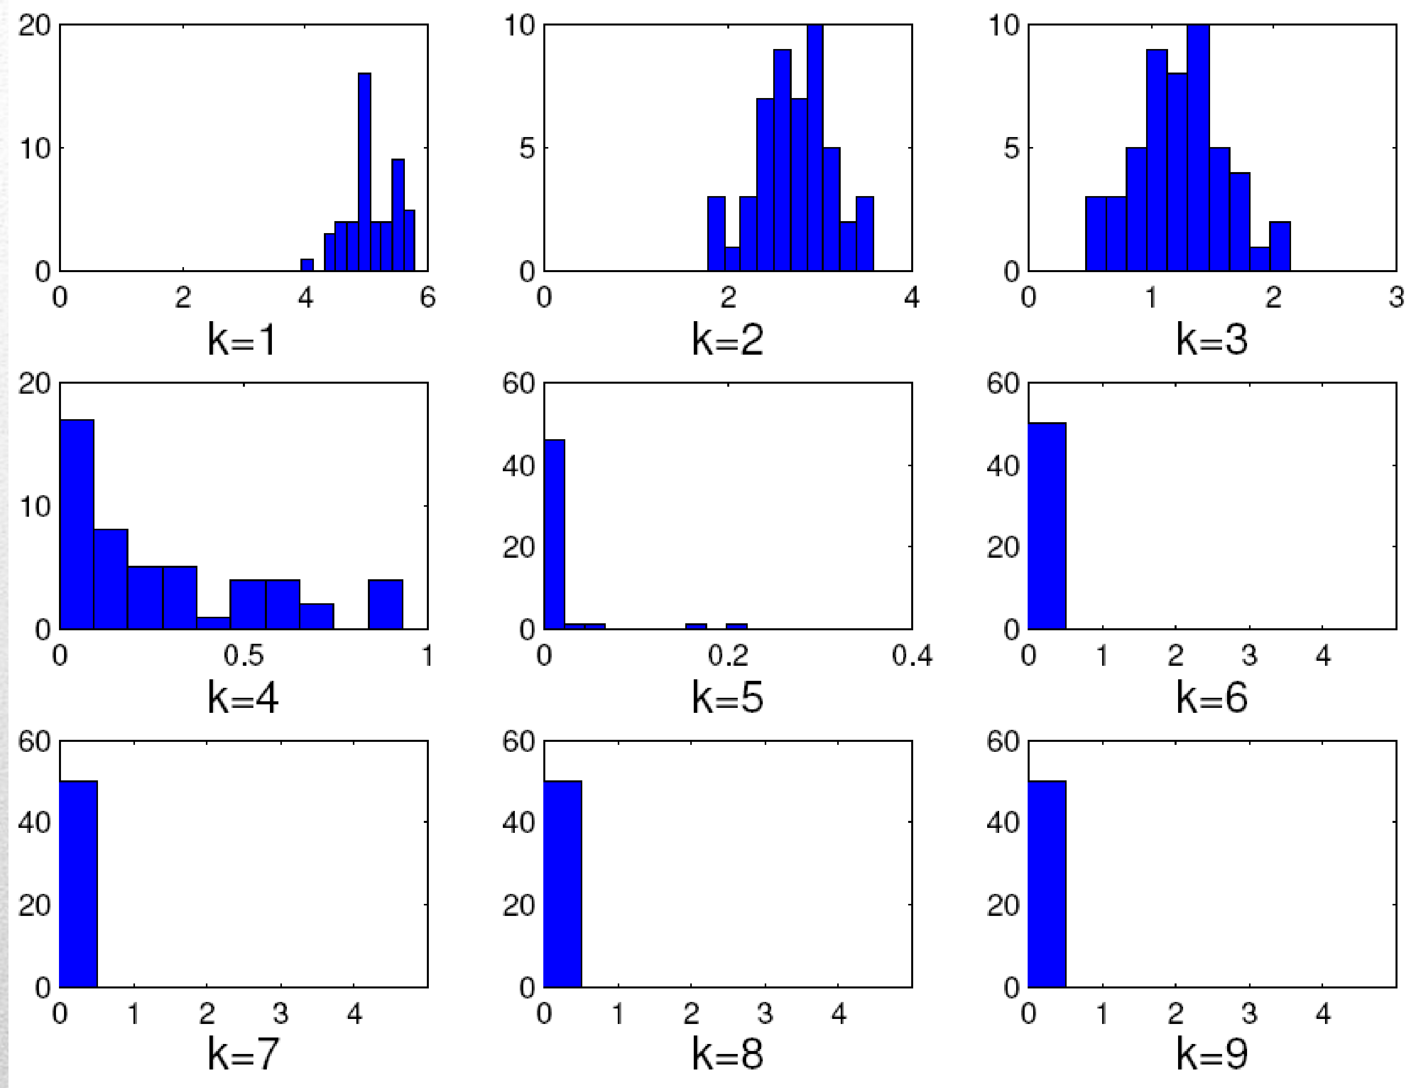

MMR Histograms: Sushi Data Set (50 samples, 100 voters each)

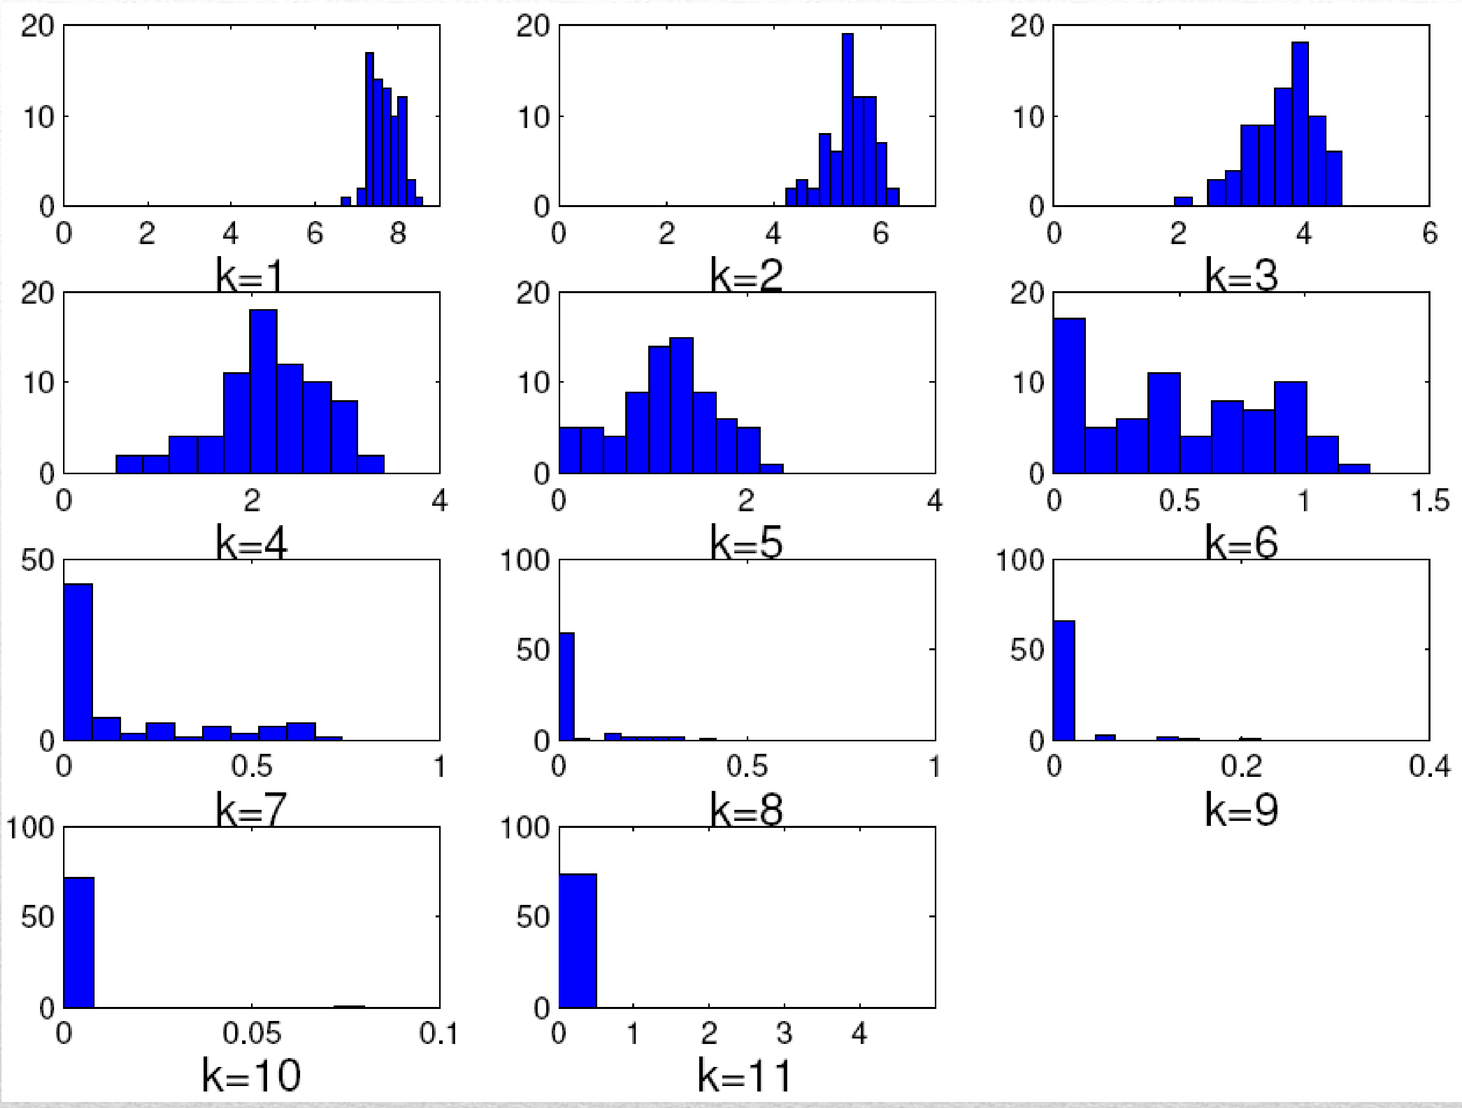

MMR Histograms: Dublin Data Set (73 samples, 50 voters each)

- Where do probabilistic models come from?
  - can be learned from sample/survey/historical data
  - two key difficulties: inference and learning
- Much research in stats, psychometrics, ML, etc.
    - but learning Mallows models with pairwise evidence ignored
- ***Inference task:*** given paired comparisons (partial vote) p_i, what is posterior over *i*'s ranking: $P(r \mid p_i ; \sigma, \phi)$
- ***Learning task:*** given partial profile $\mathbf{p} = (p_1, \ldots, p_n)$, what is max likelihood Mallows model/mixture?
  - *Solvable by EM if you can solve the inference task*

Learning Probabilistic Models [Lu, B. ICML-11]

- We adopt a sample-based approach
- Repeated Insertion Model *(DPR-04)*
  - generates samples (rankings) according to $P(r\,;\,\sigma, \phi)$
  - simple, very tractable model (cf. Young, Mallows)
- Our *Generalized Repeated Insertion Model (GRIM)*
  - generates samples (rankings) from to $P(r\,;\,\mathbf{p}, \sigma, \phi)$
  - problem intractable in general (#P-hard)
  - simple, very tractable approximations with bounds
    - works much better in practice than bounds suggest
  - procedure is exact in many important special cases
    - *E.g., samples are full rankings, top-k or partitioned preferences*

Attacking the Inference Problem*

- With sampling procedure in hand, can learn Mallows mixtures using EM from pairwise preferences
  - Tackled previously only using full voter rankings *(Murphy, Martin 2003)* or top-k *(Busse, et al. 2007)*
  - We use generalized EM with (GRIM) sample-based inference for computing expectations

Tackling the Learning Problem**

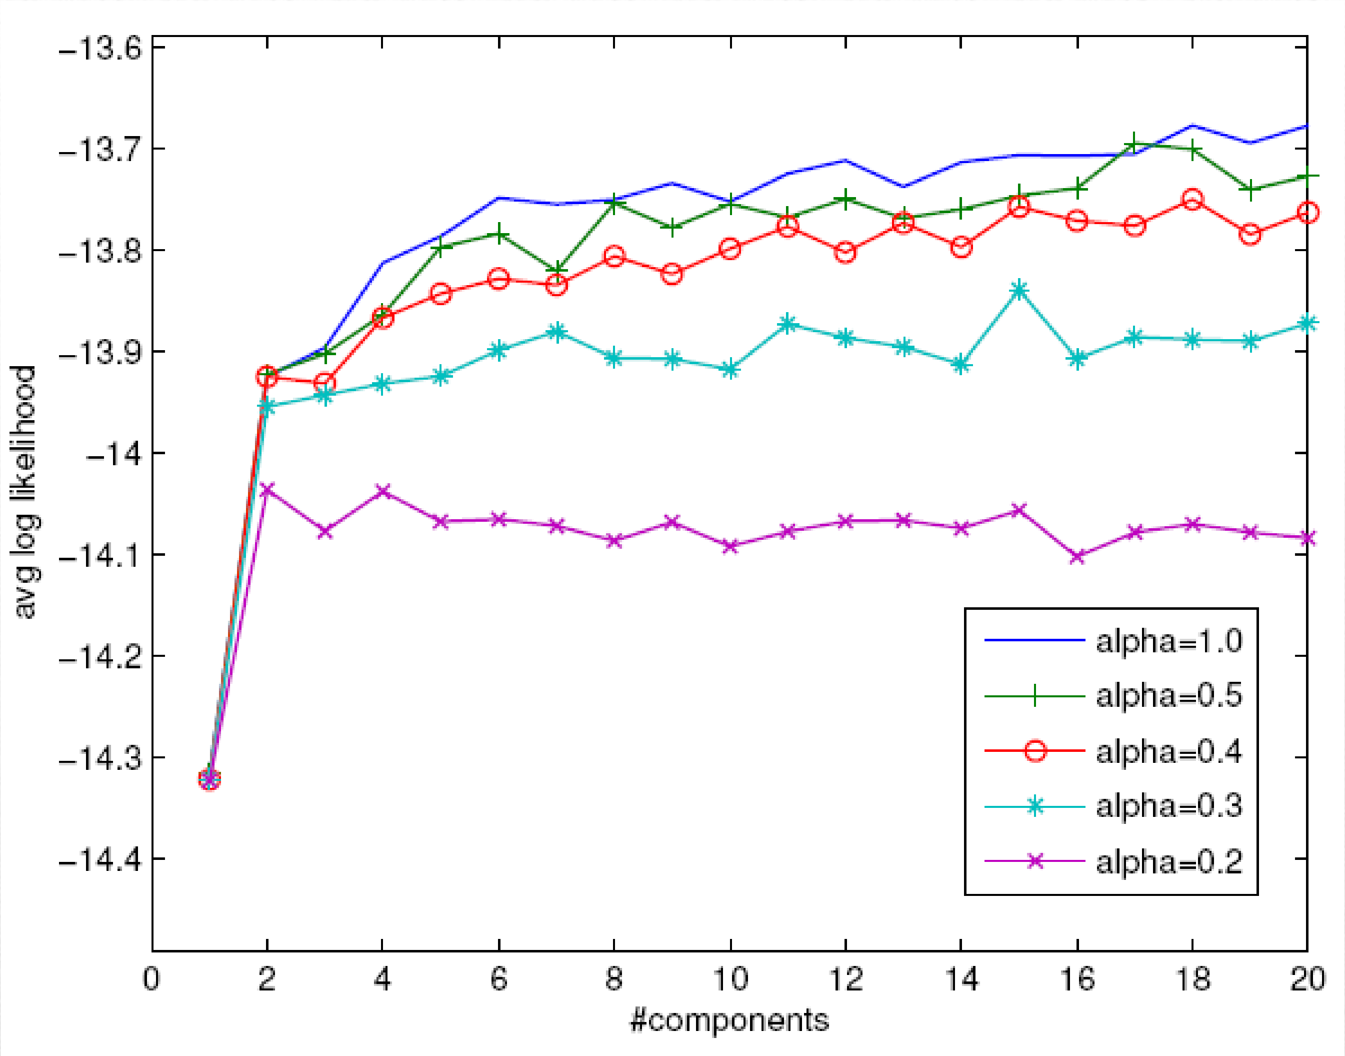

Learning Results (Sushi)**

$\pi_0 = 0.17$ $\phi_0 = 0.66$	$\pi_1 = 0.15$ $\phi_1 = 0.74$	$\pi_2 = 0.17$ $\phi_2 = 0.61$
fatty tuna	shrimp	sea urchin
salmon roe	sea eel	fatty tuna
tuna	squid	sea eel
sea eel	egg	salmon roe
tuna roll	fatty tuna	shrimp
shrimp	tuna	tuna
egg	tuna roll	squid
squid	cucumber roll	tuna roll
cucumber roll	salmon roe	egg
sea urchin	sea urchin	cucumber roll

$\pi_3 = 0.18$ $\phi_3 = 0.64$	$\pi_4 = 0.16$ $\phi_4 = 0.61$	$\pi_5 = 0.18$ $\phi_5 = 0.62$
fatty tuna	fatty tuna	fatty tuna
tuna	sea urchin	sea urchin
shrimp	tuna	salmon roe
tuna roll	salmon roe	shrimp
squid	sea eel	tuna
sea eel	tuna roll	squid
egg	shrimp	tuna roll
cucumber roll	squid	sea eel
salmon roe	egg	egg
sea urchin	cucumber roll	cucumber roll

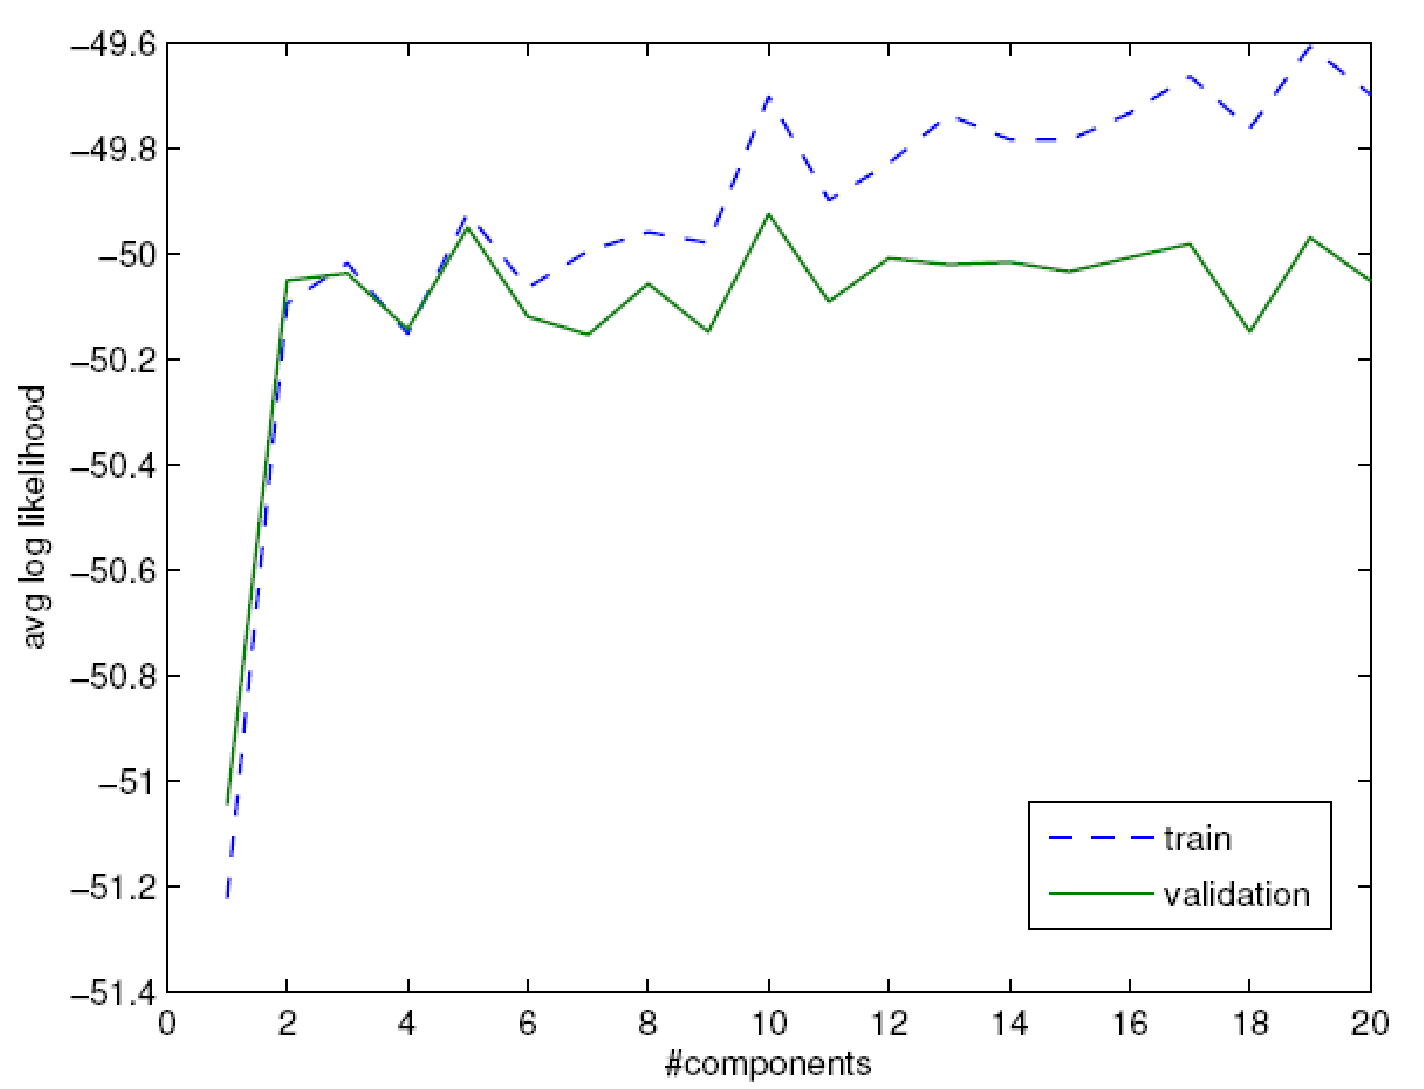

Learning Results (MovieLens)**

- Group choice: items with combinatorial structure
  - e.g., schedules, products for group use, organizational decisions (e.g., sourcing), multi-issue elections, etc.
  - representation a key issue (e.g., CP-nets)
- Minimax regret used for *single-agent* robust optimization, elicitation in combinatorial domains
  - product configuration
  - sourcing/procurement
  - resource allocation (e.g., autonomic computing)
- *Do optimization, elicitation methods extend to voting?*

Combinatorial Preference Aggregation

Computing Minimax Regret

- Example: COP with additive objective

$$\max \sum_i w_i x_i = \mathbf{w}\cdot\mathbf{x} \quad s.t.\ \mathbf{x} \in \mathbf{X}_f$$

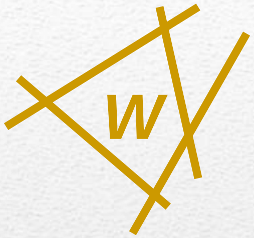

- Utility parameters ***w*** unknown: $\mathbf{w} \in \mathbf{W}$
  - difficulties computing minimax regret
  - minimax (integer) program with quadratic objective

$$MMR(W) = \min_{\mathbf{x}\in\mathbf{X}_f} \max_{\mathbf{w}\in W} \max_{\mathbf{x}'\in\mathbf{X}_f} \mathbf{w}\cdot\mathbf{x}' - \mathbf{w}\cdot\mathbf{x}$$

- General Approach:
  - Benders' decomp, constraint generation: minimax program
  - various encoding tricks to linearize quadratic terms

- Convert MMR to (linear) IP with infinitely many constraints

$$\min_{\boldsymbol{x}\in\boldsymbol{X}} \delta$$
$$\text{s.t. } \delta \geq \sum_{i=1}^{k} w_i x'_i - w_i x_i\,; \quad \forall \boldsymbol{x}' \in \boldsymbol{X}_f, \boldsymbol{w} \in \boldsymbol{W}$$

- *Repeatedly solve*

$$\min_{\boldsymbol{x}\in\boldsymbol{X}} \delta$$
$$\text{s.t. } \delta \geq \sum_{i=1}^{k} w_i x'_i - w_i x_i\,; \quad \forall (\boldsymbol{x}', \boldsymbol{w}) \in \text{Gen}$$

  - Let solution be x^* with objective value δ^*
  - Compute *MR(x*,W)* of solution x^*: *MR = r*, witness *(x'', w'')*
    - if $r > \delta^*$, add *(x'', w'')* to *Gen*, repeat; else terminate

MMR: Constraint Generation**

Regret-based Elicitation

- MMR computation effective
  - excellent anytime performance (upper, lower bounds)
- Current solution heuristic very effective for elicitation
  - typically very few queries
  - successful user studies
- Applied in several large domains
  - sourcing/combinatorial auctions
  - apartment search, product configuration
  - autonomic computing (resource alloct'n)
  - assistive technologies, etc…

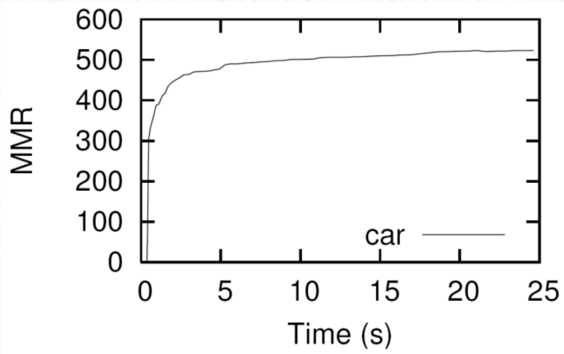

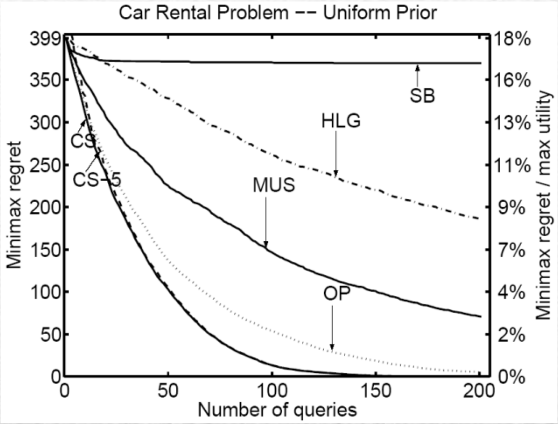

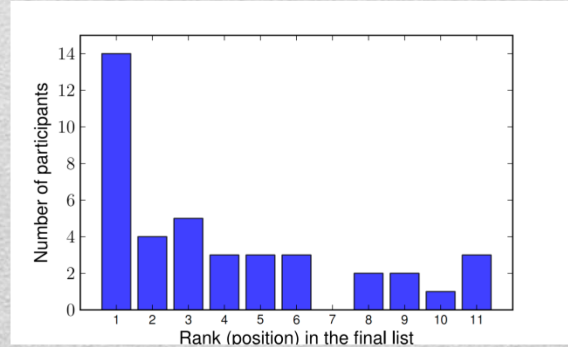

Challenges in Combinatorial Aggregation

- Tackling group elicitation in combinatorial domains
  - optimization difficult already in single agent domains
  - more subtle tradeoffs: quality, computation, elicitation burden
  - preference aggregation schemes more complex
    - CP-nets **[Rossi, Venable, Walsh; Lang, Xia, Conitzer, Maudet; Li, …];** GAI networks **[Gonzales, Perny],** etc.
    - qualitative vs. quantitative individual preferences
    - aggregate based on *total* preference? or attribute-wise?
    - if sequential, what are proper voting strategies? (equilibrium reasoning required)
    - if sequential/partial, how to optimize order?

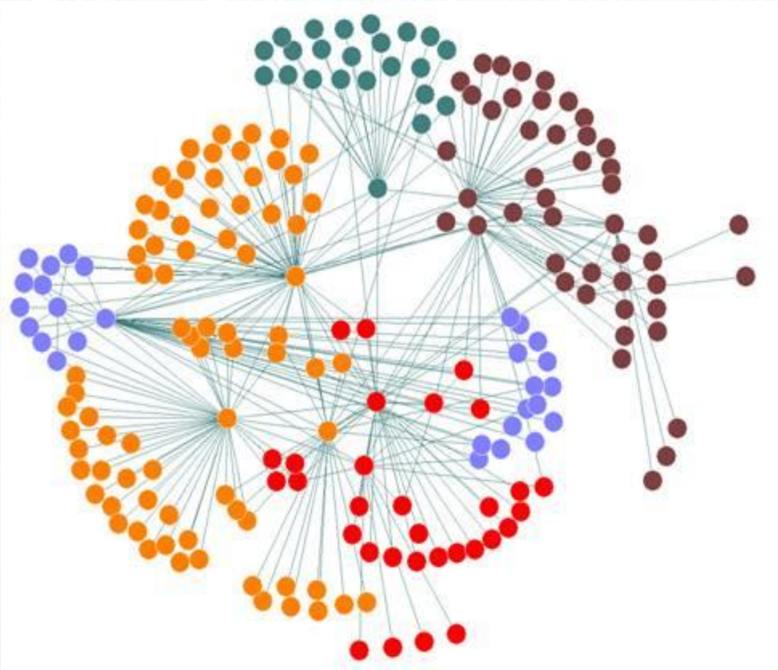

- Social networks shape behavior
  - Homophily well-documented
  - Often claimed that preferences correlated; but less evidence to this effect

- Valuable source of preference data: probabilistic models of preference correlation on networks?
  - impact on elicitation could be immense
  - both for individual or social choice problems

Social Networks as Preference Source

Social Choice on Social Networks

- Many social choice problems occur in *network context*
  - e.g., externalities in assignment (BGM EC-12), matching (BLCHW10), voting (ABKLT EC-12), coalition formation (BL11)
- Voting with *empathetic preferences* **[Saheli-Abari, B. 12]**
  - utility trades off *intrinsic* and *empathetic* preference
  - e.g., casual group decision, elections, supply chain, …
- Many new elicitation, optimization challenges

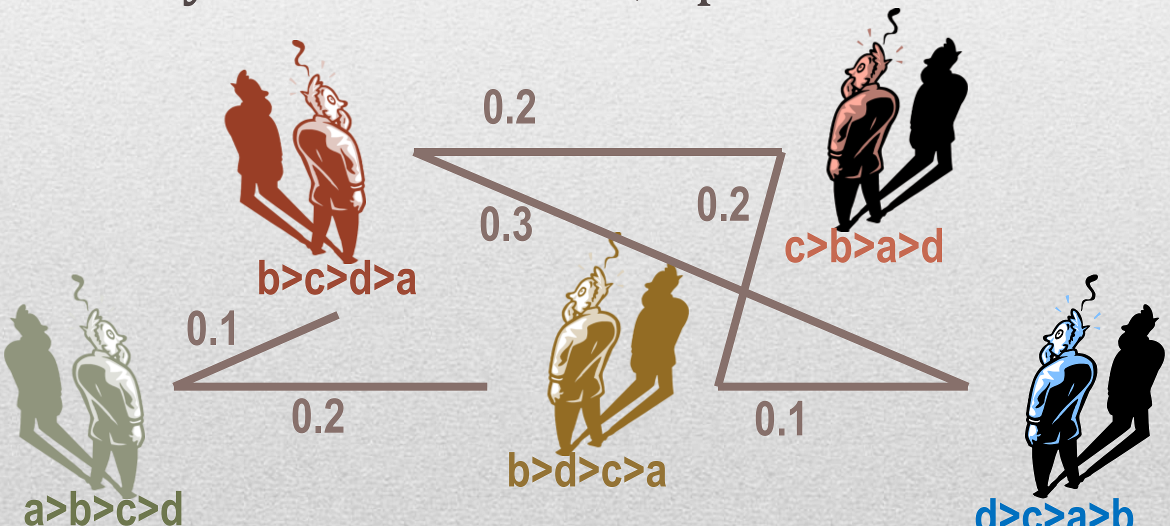

Fixed point solution (à la PageRank): Simple weighted voting scheme.

Next Steps

- Just a starting point: *learning, probabilistic models, decision-theoretic optimization* for effective elicitation and decision making in social choice settings
  - Move toward behavioural SC, connections to social media
- Next steps
  - Sophisticated, distribution-aware elicitation schemes
  - Learning other distributional models (e.g., Plackett-Luce)
  - Distributions over multi-attribute preference domains
  - Exploiting social media: networks, CF, sentiment, …
  - Computation, elicitation in combinatorial domains
  - New analyses of manipulation
  - Other social choice problems: matching; multi-winner/segmentation; allocation; etc.