Mean-Field Models and Arrow's Theorem

Frederic Koehler

INFORMS Applied Probability Society (APS) Conference, June 2025

FK is at the Department of Statistics and Data Science Institute, University of Chicago.
This talk is primarily based on a joint work with Elchanan Mossel (MIT, Department of Mathematics): "A Phase Transition in Arrow's Theorem with Three Alternatives", Annals of Applied Probability 2024.

Outline

Voting and preference aggregation

Alternatives and Preferences

- Suppose a collection of conference attendees want to collectively pick a restaurant to go to dinner at.

Alternatives and Preferences

- Suppose a collection of conference attendees want to collectively pick a restaurant to go to dinner at.
- Suppose we have a set of $q = 3$ alternatives: A, B, C.

Alternatives and Preferences

- Suppose a collection of conference attendees want to collectively pick a restaurant to go to dinner at.
- Suppose we have a set of $q = 3$ alternatives: A, B, C.
- Each voter has a *preference ranking*: a linear ordering over (A, B, C), e.g. $A \succ B \succ C$.

Alternatives and Preferences

- Suppose a collection of conference attendees want to collectively pick a restaurant to go to dinner at.
- Suppose we have a set of $q = 3$ alternatives: A, B, C.
- Each voter has a *preference ranking*: a linear ordering over (A, B, C), e.g. $A \succ B \succ C$.
- Denote by S_3 the set of all $3! = 6$ possible rankings.

Alternatives and Preferences

- Suppose a collection of conference attendees want to collectively pick a restaurant to go to dinner at.
- Suppose we have a set of $q = 3$ alternatives: A, B, C.
- Each voter has a *preference ranking*: a linear ordering over (A, B, C), e.g. $A \succ B \succ C$.
- Denote by S_3 the set of all $3! = 6$ possible rankings.
- Makes sense for general q.

Collective Decision Rules

- Key question: How can we aggregate everybody's preferences to (e.g.) pick a restaurant/president?

Collective Decision Rules

- Key question: How can we aggregate everybody's preferences to (e.g.) pick a restaurant/president?
- A *voting rule* (constitution) maps a profile of n rankings $(x_1, \ldots, x_n) \in S_3^n$ to a societal ranking in S_3.

Collective Decision Rules

- Key question: How can we aggregate everybody's preferences to (e.g.) pick a restaurant/president?
- A *voting rule* (constitution) maps a profile of n rankings $(x_1, \ldots, x_n) \in S_3^n$ to a societal ranking in S_3.
- Two desirable properties:

Collective Decision Rules

- Key question: How can we aggregate everybody's preferences to (e.g.) pick a restaurant/president?
- A *voting rule* (constitution) maps a profile of n rankings $(x_1, \ldots, x_n) \in S_3^n$ to a societal ranking in S_3.
- Two desirable properties:
  - **Unanimity**: if all voters rank $A \succ B$, then society ranks $A \succ B$.

Collective Decision Rules

- Key question: How can we aggregate everybody's preferences to (e.g.) pick a restaurant/president?
- A *voting rule* (constitution) maps a profile of n rankings $(x_1, \ldots, x_n) \in S_3^n$ to a societal ranking in S_3.
- Two desirable properties:
  - **Unanimity**: if all voters rank $A \succ B$, then society ranks $A \succ B$.
  - **Independence of Irrelevant Alternatives (IIA)**: the social preference between any two alternatives A, B depends only on how each voter ranks A vs. B.

Transitivity and Condorcet Cycles

- A societal ranking is *transitive* if $A \succ B$ and $B \succ C$ together imply $A \succ C$.

Transitivity and Condorcet Cycles

- A societal ranking is *transitive* if $A \succ B$ and $B \succ C$ together imply $A \succ C$.
- A **Condorcet winner** is an alternative that beats each other alternative in pairwise majority votes.

Transitivity and Condorcet Cycles

- A societal ranking is *transitive* if $A \succ B$ and $B \succ C$ together imply $A \succ C$.
- A **Condorcet winner** is an alternative that beats each other alternative in pairwise majority votes.
- However, pairwise majority can produce a Condorcet paradox:

Transitivity and Condorcet Cycles

- A societal ranking is *transitive* if $A \succ B$ and $B \succ C$ together imply $A \succ C$.
- A **Condorcet winner** is an alternative that beats each other alternative in pairwise majority votes.
- However, pairwise majority can produce a Condorcet paradox:
  - Voter 1: $A \succ B \succ C$

Transitivity and Condorcet Cycles

- A societal ranking is *transitive* if $A \succ B$ and $B \succ C$ together imply $A \succ C$.
- A **Condorcet winner** is an alternative that beats each other alternative in pairwise majority votes.
- However, pairwise majority can produce a Condorcet paradox:
  - Voter 1: $A \succ B \succ C$
  - Voter 2: $B \succ C \succ A$

Transitivity and Condorcet Cycles

- A societal ranking is *transitive* if $A \succ B$ and $B \succ C$ together imply $A \succ C$.
- A **Condorcet winner** is an alternative that beats each other alternative in pairwise majority votes.
- However, pairwise majority can produce a Condorcet paradox:
  - Voter 1: $A \succ B \succ C$
  - Voter 2: $B \succ C \succ A$
  - Voter 3: $C \succ A \succ B$

Transitivity and Condorcet Cycles

- A societal ranking is *transitive* if $A \succ B$ and $B \succ C$ together imply $A \succ C$.
- A **Condorcet winner** is an alternative that beats each other alternative in pairwise majority votes.
- However, pairwise majority can produce a Condorcet paradox:
  - Voter 1: $A \succ B \succ C$
  - Voter 2: $B \succ C \succ A$
  - Voter 3: $C \succ A \succ B$
- Aggregated preferences: gives $A \succ B$, $B \succ C$, and $C \succ A$!

Transitivity and Condorcet Cycles

- A societal ranking is *transitive* if $A \succ B$ and $B \succ C$ together imply $A \succ C$.
- A **Condorcet winner** is an alternative that beats each other alternative in pairwise majority votes.
- However, pairwise majority can produce a Condorcet paradox:
  - Voter 1: $A \succ B \succ C$
  - Voter 2: $B \succ C \succ A$
  - Voter 3: $C \succ A \succ B$
- Aggregated preferences: gives $A \succ B$, $B \succ C$, and $C \succ A$!
- Intransitive preferences seem deeply concerning/ **"irrational"**.

Transitivity and Condorcet Cycles

- A societal ranking is *transitive* if $A \succ B$ and $B \succ C$ together imply $A \succ C$.
- A **Condorcet winner** is an alternative that beats each other alternative in pairwise majority votes.
- However, pairwise majority can produce a Condorcet paradox:
  - Voter 1: $A \succ B \succ C$
  - Voter 2: $B \succ C \succ A$
  - Voter 3: $C \succ A \succ B$
- Aggregated preferences: gives $A \succ B$, $B \succ C$, and $C \succ A$!
- Intransitive preferences seem deeply concerning/ **"irrational"**.
- Important question: can irrational outcomes be avoided by "smarter"/different voting rules like ranked choice voting?

Arrow's Impossibility Theorem

Arrow’s Theorem (Informal)

- Important question: can irrational outcomes be avoided by “smarter” voting rules like ranked choice voting?

Arrow's Theorem (Informal)

- Important question: can irrational outcomes be avoided by "smarter" voting rules like ranked choice voting?
- Arrow showed that the answer is NO.

Arrow's Theorem (Informal)

- Important question: can irrational outcomes be avoided by "smarter" voting rules like ranked choice voting?
- Arrow showed that the answer is NO.

Arrow (1950)
For $q \geq 3$, any voting rule satisfying Unanimity and IIA must be either:

Arrow's Theorem (Informal)

- Important question: can irrational outcomes be avoided by "smarter" voting rules like ranked choice voting?
- Arrow showed that the answer is NO.

Arrow (1950)
For $q \geq 3$, any voting rule satisfying Unanimity and IIA must be either:

- a *dictatorship* (one voter determines the outcome), or

Arrow's Theorem (Informal)

- Important question: can irrational outcomes be avoided by "smarter" voting rules like ranked choice voting?
- Arrow showed that the answer is NO.

Arrow (1950)
For $q \geq 3$, any voting rule satisfying Unanimity and IIA must be either:

- a *dictatorship* (one voter determines the outcome), or
- produce an intransitive (irrational) result for some profile.

How likely are “irrational” outcomes?

Probability of a Cycle under Random Voting

- Fundamental question: Arrow's theorem tells us that intransitive outcomes are possible, but are they likely ?

Probability of a Cycle under Random Voting

- Fundamental question: Arrow's theorem tells us that intransitive outcomes are possible, but are they likely ?
- In 1952, Guilbaud started to look at this question mathematically, for the special case of majority rule/Condorcet paradox.

Probability of a Cycle under Random Voting

- Fundamental question: Arrow's theorem tells us that intransitive outcomes are possible, but are they likely ?
- In 1952, Guilbaud started to look at this question mathematically, for the special case of majority rule/Condorcet paradox.
- Assume each voter draws a ranking uniformly from S_3.

Probability of a Cycle under Random Voting

- Fundamental question: Arrow's theorem tells us that intransitive outcomes are possible, but are they likely ?
- In 1952, Guilbaud started to look at this question mathematically, for the special case of majority rule/Condorcet paradox.
- Assume each voter draws a ranking uniformly from S_3.
- As $n \to \infty$, the probability that majority rule leads to a Condorcet paradox approaches $\approx 8\%$. (Guilbaud 1952)

Probability of a Cycle under Random Voting

- Fundamental question: Arrow's theorem tells us that intransitive outcomes are possible, but are they likely ?
- In 1952, Guilbaud started to look at this question mathematically, for the special case of majority rule/Condorcet paradox.
- Assume each voter draws a ranking uniformly from S_3.
- As $n \to \infty$, the probability that majority rule leads to a Condorcet paradox approaches $\approx 8\%$. (Guilbaud 1952)
- What happens for voting rules?

Quantitative Arrow (i.i.d. voters)

Theorem (Mossel 2010, $q = 3$)
Any IIA rule that is ε-far from all dictatorships (in normalized Hamming distance) produces an intransitive outcome with probability at least $\delta(\varepsilon) > 0$, uniformly over n.

Quantitative Arrow (i.i.d. voters)

Theorem (Mossel 2010, $q = 3$)
Any IIA rule that is ε-far from all dictatorships (in normalized Hamming distance) produces an intransitive outcome with probability at least $\delta(\varepsilon) > 0$, uniformly over n.

- Formalizes a robust version: only near-dictators have low paradox risk.

Quantitative Arrow (i.i.d. voters)

Theorem (Mossel 2010, $q = 3$)
Any IIA rule that is ε-far from all dictatorships (in normalized Hamming distance) produces an intransitive outcome with probability at least $\delta(\varepsilon) > 0$, uniformly over n.

- Formalizes a robust version: only near-dictators have low paradox risk.
- Builds on line of work in discrete Fourier analysis and quantitative social choice theory (e.g., "majority is stablest", see paper for references).

Correlated preferences and statistical physics

Motivation for Correlations

- Real electorates are not independent: social influence, media, peer effects.

Motivation for Correlations

- Real electorates are not independent: social influence, media, peer effects.
- Natural to introduce correlation via Gibbs measures (statistical mechanics).

Motivation for Correlations

- Real electorates are not independent: social influence, media, peer effects.
- Natural to introduce correlation via Gibbs measures (statistical mechanics).
- Have statistical physicists thought about this problem?

Motivation for Correlations

- Real electorates are not independent: social influence, media, peer effects.
- Natural to introduce correlation via Gibbs measures (statistical mechanics).
- Have statistical physicists thought about this problem?
- YES.

Enter Statistical Physic(ist)s

Raffaelli, G. & Marsili, M. "Statistical mechanics of majority rule in large electorates," *J. Stat. Mech.* (2007).

- *Ising-like model of voting:* represented each voter by a spin encoding pairwise preferences, with ferromagnetic couplings favoring agreement.

Enter Statistical Physic(ist)s

Raffaelli, G. & Marsili, M. "Statistical mechanics of majority rule in large electorates," *J. Stat. Mech.* (2007).

- *Ising-like model of voting:* represented each voter by a spin encoding pairwise preferences, with ferromagnetic couplings favoring agreement.
- *Mean-field analysis:* derived a phase diagram showing an order–disorder transition between consensus and paradox/cycle phases as noise (inverse temperature) varies.

Enter Statistical Physic(ist)s

Raffaelli, G. & Marsili, M. "Statistical mechanics of majority rule in large electorates," *J. Stat. Mech.* (2007).

- *Ising-like model of voting:* represented each voter by a spin encoding pairwise preferences, with ferromagnetic couplings favoring agreement.
- *Mean-field analysis:* derived a phase diagram showing an order–disorder transition between consensus and paradox/cycle phases as noise (inverse temperature) varies.
- *Critical behaviour:* identified the critical inverse temperature β_c at which the probability of a Condorcet paradox drops to zero, and characterized fluctuations near β_c. (Did not consider general voting rules.)

Enter Statistical Physic(ist)s

Raffaelli, G. & Marsili, M. "Statistical mechanics of majority rule in large electorates," *J. Stat. Mech.* (2007).

- *Ising-like model of voting:* represented each voter by a spin encoding pairwise preferences, with ferromagnetic couplings favoring agreement.
- *Mean-field analysis:* derived a phase diagram showing an order–disorder transition between consensus and paradox/cycle phases as noise (inverse temperature) varies.
- *Critical behaviour:* identified the critical inverse temperature β_c at which the probability of a Condorcet paradox drops to zero, and characterized fluctuations near β_c. (Did not consider general voting rules.)
- *Challenge*: their analysis is not a (mathematically) rigorous proof. See open problems later !

Mean-field (Rafaelli-Marsili) model

Mean-field model, Rafaelli-Marsili '07

Preferences $X = (X_i)_{i=1}^n \in S_q^n$ with energy

$$E(X) = \sum_{1 \le i < j \le n} 2 d_\tau(X_i, X_j),$$

where d_τ is Kendall's tau distance between rankings.

Gibbs distribution at inverse temperature β:

$$Q_\beta(X) = \frac{1}{Z} \exp\Big(-\frac{2\beta}{n} E(X)\Big).$$

Captures mean-field interaction among all pairs.

Large Deviations Interpretation

Define $\varphi : \mathfrak{S}_q \to \{\pm 1\}^{\binom{q}{2}}$ by

$$\varphi(\pi)_{i,j} = (-1)^{1[\pi(i) > \pi(j)]}$$

Then Kendall's tau distance has a "kernelized" representation:

$$\langle \varphi(\pi), \varphi(\pi') \rangle = \binom{q}{2} - 2d_\tau(\pi, \pi')$$

By Cramer's theorem, studying the mean-field model is essentially equivalent to understanding *large deviations* of i.i.d. permutations: i.e. for sets $S \in \mathbb{R}^{\binom{q}{2}}$, computing

$$\log \Pr\left[\frac{1}{n}\sum_{i=1}^{n} \varphi(\pi_i) \in S\right]$$

where $\pi_1, \ldots, \pi_n \sim \text{Uni}\mathfrak{S}_q$. See [Ellis, '06]

Main Results (K-Mossel '24)

High-Temperature Phase ($\beta < 3/4$)

Our main results rigorously solve the statistical mechanics model and show that its phase transition completely determines the behavior of Arrow's Theorem when $q = 3$.

Theorem (Quantitative Arrow, Mean-Field, $q = 3$)
Fix $\beta < 3/4$ and $\varepsilon > 0$. Any IIA rule ε-far from dictatorships has paradox probability at least $\delta(\varepsilon, \beta) > 0$, uniformly in n.

High-Temperature Phase ($\beta < 3/4$)

Our main results rigorously solve the statistical mechanics model and show that its phase transition completely determines the behavior of Arrow's Theorem when $q = 3$.

Theorem (Quantitative Arrow, Mean-Field, $q = 3$)
Fix $\beta < 3/4$ and $\varepsilon > 0$. Any IIA rule ε-far from dictatorships has paradox probability at least $\delta(\varepsilon, \beta) > 0$, uniformly in n.

- Paradox remains nontrivial in correlated high-T regime.

High-Temperature Phase ($\beta < 3/4$)

Our main results rigorously solve the statistical mechanics model and show that its phase transition completely determines the behavior of Arrow's Theorem when $q = 3$.

Theorem (Quantitative Arrow, Mean-Field, $q = 3$)
Fix $\beta < 3/4$ and $\varepsilon > 0$. Any IIA rule ε-far from dictatorships has paradox probability at least $\delta(\varepsilon, \beta) > 0$, uniformly in n.

- Paradox remains nontrivial in correlated high-T regime.
- Threshold of $3/4$ agrees with physics analysis for majority rule $+$ order-disorder phase transition.

High-Temperature Phase ($\beta < 3/4$)

Our main results rigorously solve the statistical mechanics model and show that its phase transition completely determines the behavior of Arrow's Theorem when $q = 3$.

Theorem (Quantitative Arrow, Mean-Field, $q = 3$)
Fix $\beta < 3/4$ and $\varepsilon > 0$. Any IIA rule ε-far from dictatorships has paradox probability at least $\delta(\varepsilon, \beta) > 0$, uniformly in n.

- Paradox remains nontrivial in correlated high-T regime.
- Threshold of $3/4$ agrees with physics analysis for majority rule + order-disorder phase transition.
- Established via contiguity.

Phase Transition in Majority Rule

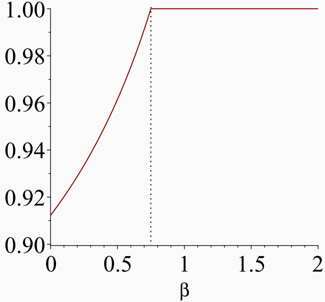

For majority voting, the no-cycle probability converges to

$$p_\infty(\beta) = \begin{cases} \frac{3}{2\pi} \arccos\left(\frac{3-4\beta}{9}\right), & \beta < \frac{3}{4}, \\ 1, & \beta > \frac{3}{4}. \end{cases}$$

- Shows a sharp threshold at $\beta_c = 3/4$.

Phase Transition in Majority Rule

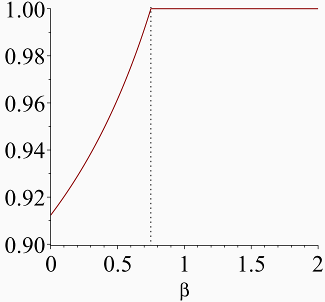

For majority voting, the no-cycle probability converges to

$$p_\infty(\beta) = \begin{cases} \frac{3}{2\pi} \arccos\left(\frac{3-4\beta}{9}\right), & \beta < \frac{3}{4}, \\ 1, & \beta > \frac{3}{4}. \end{cases}$$

- Shows a sharp threshold at $\beta_c = 3/4$.
- Below β_c: bounded away from 1; above: paradox vanishes.

Low-Temperature Behavior ($\beta > 3/4$)

- Theorem: For $\beta > 3/4$, majority rule paradox probability decays exponentially in n.

Low-Temperature Behavior ($\beta > 3/4$)

- Theorem: For $\beta > 3/4$, majority rule paradox probability decays exponentially in n.
- So Quantitative Arrow fails: near-majority functions avoid cycles.

Low-Temperature Behavior ($\beta > 3/4$)

- Theorem: For $\beta > 3/4$, majority rule paradox probability decays exponentially in *n*.
- So Quantitative Arrow fails: near-majority functions avoid cycles.
- Analysis is connected to physics ideas (mean-field approximation).

Low-Temperature Behavior ($\beta > 3/4$)

- Theorem: For $\beta > 3/4$, majority rule paradox probability decays exponentially in n.
- So Quantitative Arrow fails: near-majority functions avoid cycles.
- Analysis is connected to physics ideas (mean-field approximation).
- **Rigorous mean-field approximation to free energy:**

$$\frac{1}{n} \log Z_n \to \max_{s \in [-1,1]^3} \Phi(s).$$

where $NAE_3 = \{\pm 1\}^3 \setminus \{(1,1,1),(-1,-1,-1)\}$ and

$$\Phi(s) := \frac{\beta}{2} \|s\|_2^2 + \max_{Q \in \mathcal{P}(NAE_3): \mathbb{E}_Q X = s} H_Q(X).$$

Low-Temperature Behavior ($\beta > 3/4$)

- Theorem: For $\beta > 3/4$, majority rule paradox probability decays exponentially in n.
- So Quantitative Arrow fails: near-majority functions avoid cycles.
- Analysis is connected to physics ideas (mean-field approximation).
- **Rigorous mean-field approximation to free energy:**

$$\frac{1}{n} \log Z_n \to \max_{s \in [-1,1]^3} \Phi(s).$$

where $NAE_3 = \{\pm 1\}^3 \setminus \{(1,1,1),(-1,-1,-1)\}$ and

$$\Phi(s) := \frac{\beta}{2} \|s\|_2^2 + \max_{Q \in \mathcal{P}(NAE_3) : \mathbb{E}_Q X = s} H_Q(X).$$

- Key mathematical difficulty: analyzing the limiting variational problem is not easy !

Mean-field equations

Define for $s \in [-1,1]^3$ the functional

$$\Phi(s) := \frac{\beta}{2}\|s\|_2^2 + \max_{Q \in \mathcal{P}(NAE_3): \mathbb{E}_Q X = s} H_Q(X).$$

Then at any critical point of Φ, we have the equality

$$\Phi(s) = \log(\cosh(\lambda_1+\lambda_2-\lambda_3) + \cosh(\lambda_1-\lambda_2+\lambda_3) + \cosh(-\lambda_1+\lambda_2+\lambda_3)) - \frac{\beta}{2}\|s\|_2^2 + \log 2$$

where $\lambda_i = \beta s_i$ for $i = 1,2,3$ which satisfy *mean-field equations*

$$s_1 = \frac{\sinh(\lambda_1+\lambda_2-\lambda_3) + \sinh(\lambda_1-\lambda_2+\lambda_3) - \sinh(-\lambda_1+\lambda_2+\lambda_3)}{\cosh(\lambda_1+\lambda_2-\lambda_3) + \cosh(\lambda_1-\lambda_2+\lambda_3) + \cosh(-\lambda_1+\lambda_2+\lambda_3)}$$

$$s_2 = \frac{\sinh(\lambda_1+\lambda_2-\lambda_3) - \sinh(\lambda_1-\lambda_2+\lambda_3) + \sinh(-\lambda_1+\lambda_2+\lambda_3)}{\cosh(\lambda_1+\lambda_2-\lambda_3) + \cosh(\lambda_1-\lambda_2+\lambda_3) + \cosh(-\lambda_1+\lambda_2+\lambda_3)}$$

$$s_3 = \frac{-\sinh(\lambda_1+\lambda_2-\lambda_3) + \sinh(\lambda_1-\lambda_2+\lambda_3) + \sinh(-\lambda_1+\lambda_2+\lambda_3)}{\cosh(\lambda_1+\lambda_2-\lambda_3) + \cosh(\lambda_1-\lambda_2+\lambda_3) + \cosh(-\lambda_1+\lambda_2+\lambda_3)}.$$

Analysis of solutions and Theorem

KEY LEMMA: For all $\beta \geq 0$, the solutions to the mean-field equations are of one of the following types, up to symmetries of permuting coordinates and $\lambda \mapsto -\lambda$:

1. Of the form $\lambda_1 = \lambda_2 = \lambda_3$

Furthermore, points of third type are the global maximizers of Φ.

Analysis of solutions and Theorem

KEY LEMMA: For all $\beta \geq 0$, the solutions to the mean-field equations are of one of the following types, up to symmetries of permuting coordinates and $\lambda \mapsto -\lambda$:

1. Of the form $\lambda_1 = \lambda_2 = \lambda_3$
2. Of the form $\lambda_1 = 0, \lambda_2 = -\lambda_3$.

Furthermore, points of third type are the global maximizers of Φ.

Analysis of solutions and Theorem

KEY LEMMA: For all $\beta \geq 0$, the solutions to the mean-field equations are of one of the following types, up to symmetries of permuting coordinates and $\lambda \mapsto -\lambda$:

1. Of the form $\lambda_1 = \lambda_2 = \lambda_3$
2. Of the form $\lambda_1 = 0, \lambda_2 = -\lambda_3$.
3. Of the form $\lambda_1 = \lambda_2$ where λ_3 has the opposite sign of λ_1 and up to symmetries, this point is unique (for $\beta > 3/4$ it has an orbit of size exactly 6).

Furthermore, points of third type are the global maximizers of Φ.

Analysis of solutions and Theorem

KEY LEMMA: For all $\beta \geq 0$, the solutions to the mean-field equations are of one of the following types, up to symmetries of permuting coordinates and $\lambda \mapsto -\lambda$:

1. Of the form $\lambda_1 = \lambda_2 = \lambda_3$
2. Of the form $\lambda_1 = 0, \lambda_2 = -\lambda_3$.
3. Of the form $\lambda_1 = \lambda_2$ where λ_3 has the opposite sign of λ_1 and up to symmetries, this point is unique (for $\beta > 3/4$ it has an orbit of size exactly 6).

Furthermore, points of third type are the global maximizers of Φ.

Analysis of solutions and Theorem

KEY LEMMA: For all $\beta \geq 0$, the solutions to the mean-field equations are of one of the following types, up to symmetries of permuting coordinates and $\lambda \mapsto -\lambda$:

1. Of the form $\lambda_1 = \lambda_2 = \lambda_3$
2. Of the form $\lambda_1 = 0, \lambda_2 = -\lambda_3$.
3. Of the form $\lambda_1 = \lambda_2$ where λ_3 has the opposite sign of λ_1 and up to symmetries, this point is unique (for $\beta > 3/4$ it has an orbit of size exactly 6).

Furthermore, points of third type are the global maximizers of Φ.
FINAL THEOREM:

$$\frac{1}{n} S_n \to S \sim Uni\{\text{global maximizers of } \Phi\}.$$

Solution Locii (over all β)

$$u := \begin{bmatrix} 1 & 1 & -1 \\ 1 & -1 & 1 \\ -1 & 1 & 1 \end{bmatrix} \lambda.$$

Solution Locii (over all β)

$$u := \begin{bmatrix} 1 & 1 & -1 \\ 1 & -1 & 1 \\ -1 & 1 & 1 \end{bmatrix} \lambda.$$

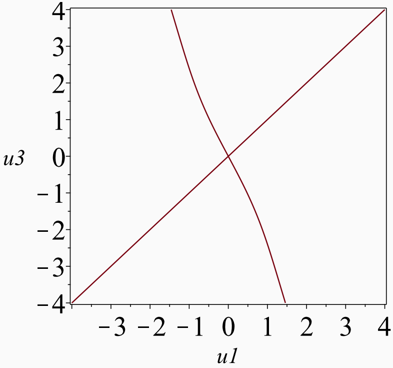

Figure 1: LEFT: solutions when $u_1 = u_2$. RIGHT: $u_1 \neq u_2$.

Solution Locii (over all β)

$$u := \begin{bmatrix} 1 & 1 & -1 \\ 1 & -1 & 1 \\ -1 & 1 & 1 \end{bmatrix} \lambda.$$

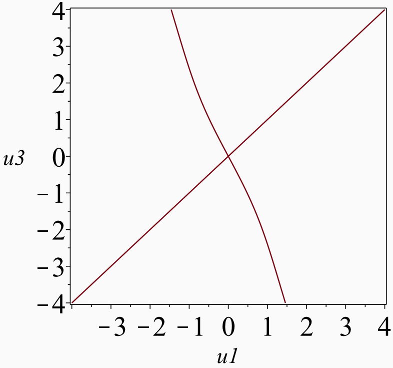

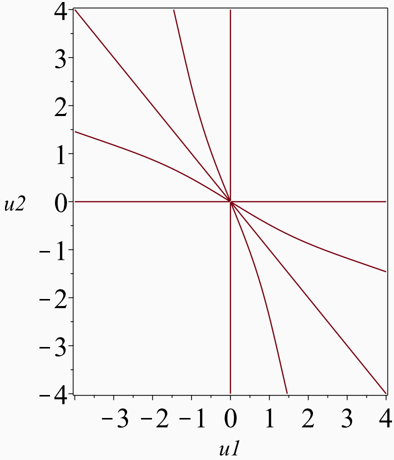

Figure 1: LEFT: solutions when $u_1 = u_2$. RIGHT: $u_1 \neq u_2$.

Beyond the mean-field model

Perfect Matching Model

Divide voters into $n/2$ disjoint pairs (i, j). Preferences sampled via

$$Q_\beta(X) \propto \exp\Big(-2\beta \sum_{(i,j)} d_\tau(X_i, X_j)\Big).$$

- We prove quantitative Arrow's Theorem holds for all β (based on generalizing ideas from i.i.d. case — tricky).

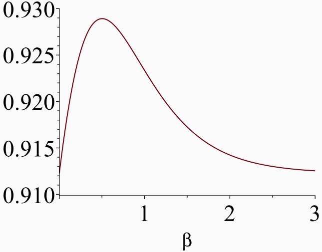

Perfect Matching Model

Divide voters into $n/2$ disjoint pairs (i, j). Preferences sampled via

$$Q_\beta(X) \propto \exp\Big(-2\beta \sum_{(i,j)} d_\tau(X_i, X_j)\Big).$$

- We prove quantitative Arrow's Theorem holds for all β (based on generalizing ideas from i.i.d. case — tricky).
- Paradox probability can be non-monotone in β. (Below: majority rule)

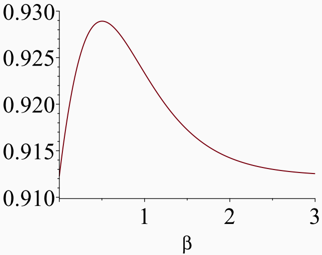

Conclusions and Open Problems

Quantitative Arrow for $q > 3$?

Theorem (Mean-Field, general q)
If $\beta < 1/(q-1)$ and a rule is ε-far from dictators, paradox probability $\geq \delta(\varepsilon, \beta) > 0$.

Quantitative Arrow for $q > 3$?

Theorem (Mean-Field, general q)
If $\beta < 1/(q-1)$ and a rule is ε-far from dictators, paradox probability $\geq \delta(\varepsilon, \beta) > 0$.

- Above another threshold ($\beta > 3/(q+1)$) we can prove the model enters a low-T regime.

Quantitative Arrow for $q > 3$?

Theorem (Mean-Field, general q)
If $\beta < 1/(q-1)$ and a rule is ε-far from dictators, paradox probability $\geq \delta(\varepsilon, \beta) > 0$.

- Above another threshold ($\beta > 3/(q+1)$) we can prove the model enters a low-T regime.
- Physics analysis (Rafaelli-Marsili): the correct answer should be a sharp "second-order" phase transition at $\beta_c = 3/(q+1)$. (Not first-order like Potts model !)

Quantitative Arrow for $q > 3$?

Theorem (Mean-Field, general q)
If $\beta < 1/(q-1)$ and a rule is ε-far from dictators, paradox probability $\geq \delta(\varepsilon, \beta) > 0$.

- Above another threshold ($\beta > 3/(q+1)$) we can prove the model enters a low-T regime.
- Physics analysis (Rafaelli-Marsili): the correct answer should be a sharp "second-order" phase transition at $\beta_c = 3/(q+1)$. (Not first-order like Potts model !)
- Predicted behavior corresponds to an interesting prediction about large deviations of random permutations.

Large Deviations Conjecture

Conjecture:

$$\log \mathbb{E}[\exp(\langle \lambda, X \rangle)] \leq \frac{q+1}{6} \|\lambda\|^2.$$

This is a sharp "sub-Gaussian" concentration inequality about weighted inversions of a random permutation.

If true, this constant is sharp, it corresponds to the predicted phase transition at $3/(q+1)$. We can prove the result with a slightly worse constant $(q-1)/2$ via a martingale argument over a (symmetrized) Fisher-Yates shuffle.

Conclusions

- Introduced Gibbs models for correlated preferences.

Conclusions

- Introduced Gibbs models for correlated preferences.
- Established a sharp phase transition at $\beta_c = 3/4$ for $q = 3$.

Conclusions

- Introduced Gibbs models for correlated preferences.
- Established a sharp phase transition at $\beta_c = 3/4$ for $q = 3$.
- Proved quantitative Arrow in high-T regime.

Conclusions

- Introduced Gibbs models for correlated preferences.
- Established a sharp phase transition at $\beta_c = 3/4$ for $q = 3$.
- Proved quantitative Arrow in high-T regime.
- Observed non-monotone paradox behavior in sparse models.

Open Problems

- Precise analysis around critical point $\beta = 3/4$.

Open Problems

- Precise analysis around critical point $\beta = 3/4$.
- Tighten thresholds and constants for general q.

Open Problems

- Precise analysis around critical point $\beta = 3/4$.
- Tighten thresholds and constants for general q.
- Many other open problems stemming from Rafaelli-Marsili's work:

Open Problems

- Precise analysis around critical point $\beta = 3/4$.
- Tighten thresholds and constants for general q.
- Many other open problems stemming from Rafaelli-Marsili's work:
    - They also analyzed behavior of model under random fields.

Open Problems

- Precise analysis around critical point $\beta = 3/4$.
- Tighten thresholds and constants for general q.
- Many other open problems stemming from Rafaelli-Marsili's work:
  - They also analyzed behavior of model under random fields.
  - Rigorously understanding low temperature/"ordered" phase for $q > 3$ is also quite complex (IMO).

Open Problems

- Precise analysis around critical point $\beta = 3/4$.
- Tighten thresholds and constants for general q.
- Many other open problems stemming from Rafaelli-Marsili's work:
  - They also analyzed behavior of model under random fields.
  - Rigorously understanding low temperature/"ordered" phase for $q > 3$ is also quite complex (IMO).
- Extend to other interaction graphs (e.g., lattices, networks). Conjecture: quantitative Arrow's theorem holds under natural Dobrushin uniqueness condition.

Open Problems

- Precise analysis around critical point $\beta = 3/4$.
- Tighten thresholds and constants for general q.
- Many other open problems stemming from Rafaelli-Marsili's work:
    - They also analyzed behavior of model under random fields.
    - Rigorously understanding low temperature/"ordered" phase for $q > 3$ is also quite complex (IMO).
- Extend to other interaction graphs (e.g., lattices, networks). Conjecture: quantitative Arrow's theorem holds under natural Dobrushin uniqueness condition.
    - Relevant recent work: sharp Kahn-Kalai-Linial (KKL) Theorem for a large class of sparse graphical models via Glauber semigroup + diagrammatic commutator expansion [K-Lifshitz-Minzer-Mossel '24]

Open Problems

- Precise analysis around critical point $\beta = 3/4$.
- Tighten thresholds and constants for general q.
- Many other open problems stemming from Rafaelli-Marsili's work:
  - They also analyzed behavior of model under random fields.
  - Rigorously understanding low temperature/"ordered" phase for $q > 3$ is also quite complex (IMO).
- Extend to other interaction graphs (e.g., lattices, networks). Conjecture: quantitative Arrow's theorem holds under natural Dobrushin uniqueness condition.
  - Relevant recent work: sharp Kahn-Kalai-Linial (KKL) Theorem for a large class of sparse graphical models via Glauber semigroup + diagrammatic commutator expansion [K-Lifshitz-Minzer-Mossel '24]
  - Removing sparsity assumption for KKL is a nice open problem.

Open Problems

- Precise analysis around critical point $\beta = 3/4$.
- Tighten thresholds and constants for general q.
- Many other open problems stemming from Rafaelli-Marsili's work:
  - They also analyzed behavior of model under random fields.
  - Rigorously understanding low temperature/"ordered" phase for $q > 3$ is also quite complex (IMO).
- Extend to other interaction graphs (e.g., lattices, networks). Conjecture: quantitative Arrow's theorem holds under natural Dobrushin uniqueness condition.
  - Relevant recent work: sharp Kahn-Kalai-Linial (KKL) Theorem for a large class of sparse graphical models via Glauber semigroup + diagrammatic commutator expansion [K-Lifshitz-Minzer-Mossel '24]
  - Removing sparsity assumption for KKL is a nice open problem.
  - Arrow's theorem seems trickier: cf. invariance principle.