Stable matchings beyond worst-case

Claire Mathieu

CNRS dépasser les frontières

Joint work with Hugo Gimbert and Simon Mauras

the Gale-Shapley deferred acceptance algorithm

Matching Universities and Candidates

Universities have preferences over candidates
Candidates have preferences over universities

Gale-Shapley algorithm
Repeat

i. each university U sends proposals to its C favorite candidates, with C the number of slots still available.
ii. each candidate with at least one proposal accepts (provisionally) her favorite proposal and declines the other proposals.
iii. each university removes from its list the candidates that have declined.

U: 1 2 3 4 5 6
V: 1 3 2 4 5 6
W: 1 3 2 4 5 6

Matching Universities and Candidates

Gale-Shapley algorithm. Assume each university has capacity 1.

Repeat

i. each university call its favorite candidate.

ii. each candidate accepts her favorite proposal and declines the other proposals.

iii. each university removes from its list the candidates that have declined.

Step1

U: 1 2 3 4 5 6
V: 1 3 2 4 5 6
W: 1 3 2 4 5 6

Step2

U: 1 2 3 4 5 6
V: 1 3 2 4 5 6
W: 1 3 2 4 5 6

Step3

U: 1 2 3 4 5 6
V: 1 3 2 4 5 6
W: 1 3 2 4 5 6

Step4

U: 1 2 3 4 5 6
V: 1 3 2 4 5 6
W: 1 3 2 4 5 6

Step5

U: 1 2 3 4 5 6
V: 1 3 2 4 5 6
W: 1 3 2 4 5 6

Step6

U: 1 2 3 4 5 6
V: 1 3 2 4 5 6
W: 1 3 2 4 5 6

Stable matching

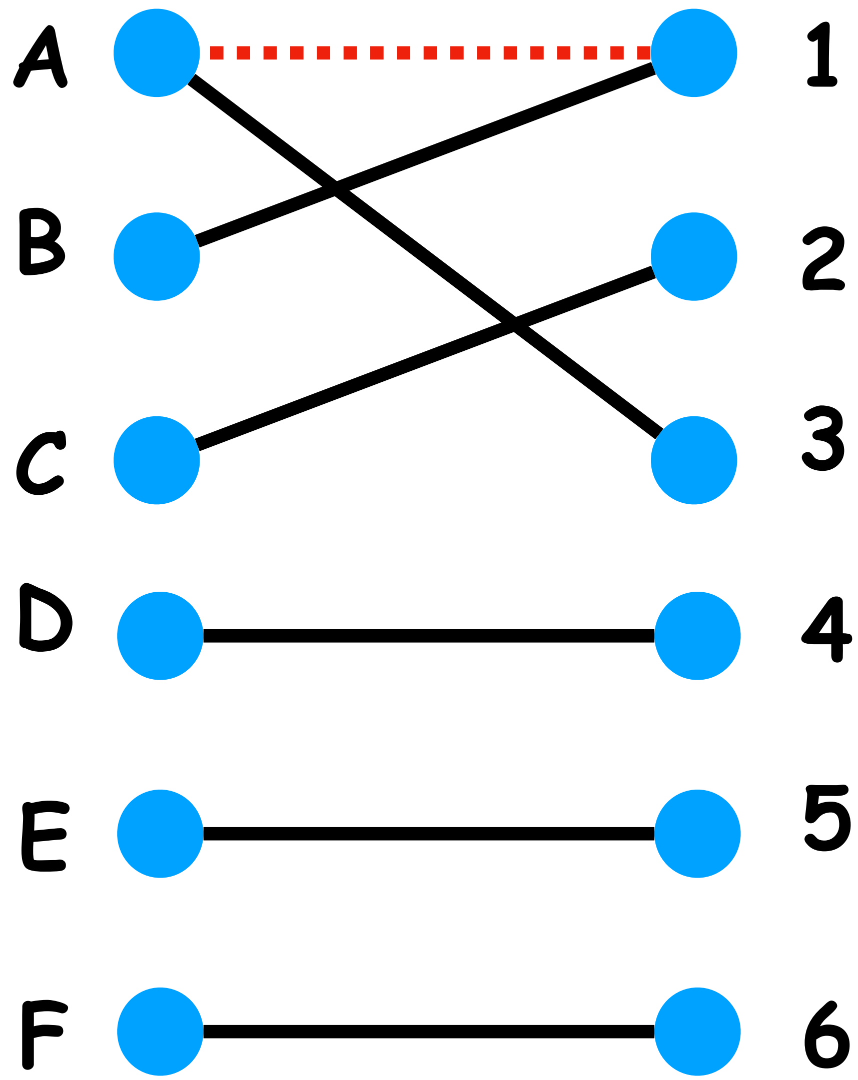

A: 1 2 3 4 5 6 none
B: 5 3 6 1 4 2 none
...
F: 6 5 1 2 3 4 none
1: A B C D E F none
2: D F E A B C none
...
6: A B C D E F none

Blocking pair:
A and 1 prefer each other to their mates

A blocking pair makes the bipartite matching **unstable**
It is desirable that a matching be **stable**

Manipulability of stable matchings

Step6						
U:	1	2	3	4	5	6
V:	1	4	3	2	5	6
W:	1	3	2	4	5	6
X:	4	1	3	2	5	6

Q. Can candidate 4 obtain a better university than X?

A. To know her options, 4 rejects all offers.

Manipulability of stable matchings

To know her options, 4 rejects all offers

Step6
U: 1 2 3 4 5 6
V: 1 4 3 2 5 6
W: 1 3 2 4 5 6
X: 4 1 3 2 5 6
4:{X}

Step7
U: 1 2 3 4 5 6
V: 1 4 2 4 5 6
W: 1 3 2 4 5 6
X: 4 1 3 2 5 6
4:{X}

Step8
U: 1 2 3 4 5 6
V: 1 4 2 4 5 6
W: 1 3 2 4 5 6
X: 4 1 3 2 5 6
4:{X,V}

Step9
U: 1 2 3 4 5 6
V: 1 4 2 4 5 6
W: 1 3 2 4 5 6
X: 4 1 3 2 5 6
4:{X,V}

Step10
U: 1 2 3 4 5 6
V: 1 4 2 4 5 6
W: 1 3 2 4 5 6
X: 4 1 3 2 5 6
4:{X,V,W}

Step11
U: 1 2 3 4 5 6
V: 1 4 2 4 5 6
W: 1 3 2 4 5 6
X: 4 1 3 2 5 6
4:{X,V,W}

Step12
U: 1 2 3 4 5 6
V: 1 4 2 4 5 6
W: 1 3 2 4 5 6
X: 4 1 3 2 5 6
4:{X,V,W,none}

Strategic manipulation: if Candidate 4 has perfect information, she can pick her favorite in {X,V,W,none} and the matching is still stable

The lattice structure of stable matchings

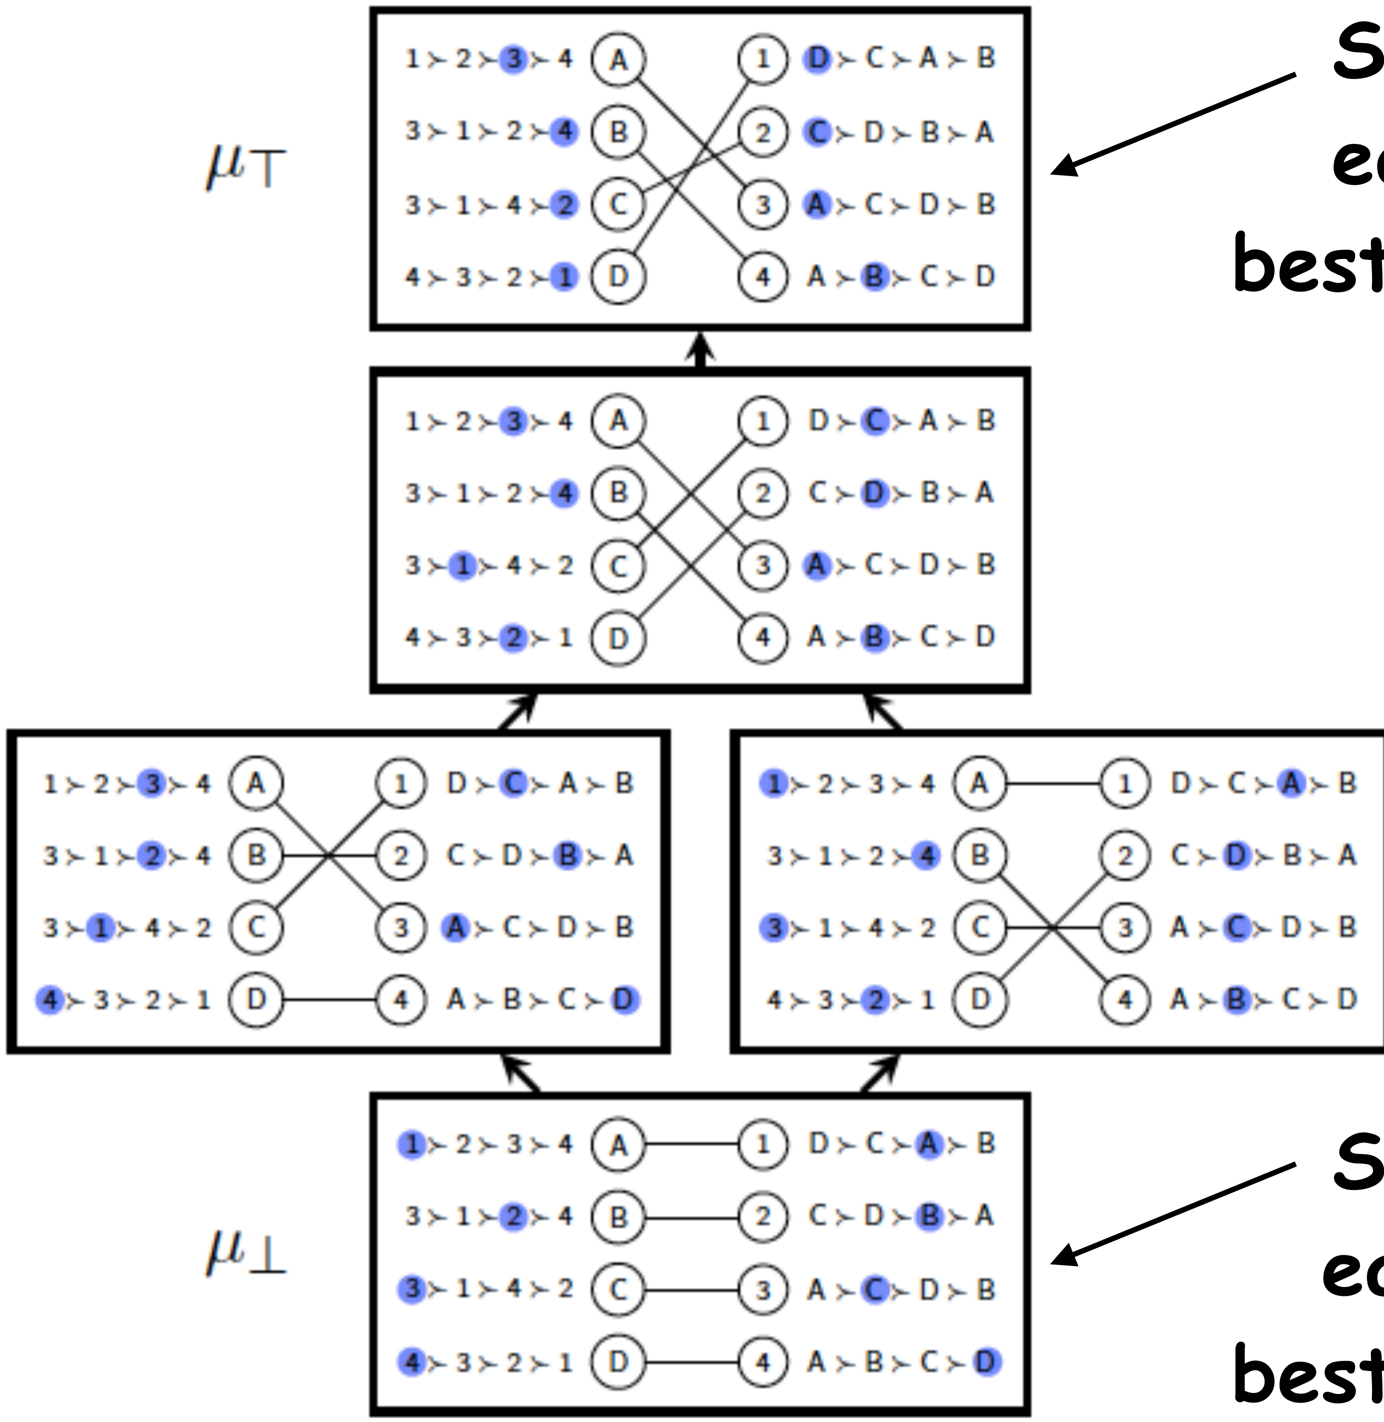

Stable matching where each university has his best possible stable partner

Stable matching where each candidate has her best possible stable partner

Non-Manipulability of Gale-Shapley…

Assume universities have capacity 1

Gale-Shapley man-optimal algorithm:
produces a stable matching
s.t. every university obtains his best possible stable partner.

No university may obtain a better outcome by lying about its preferences

… Manipulability of Gale-Shapley

Candidates may obtain a better outcome by lying about their preferences

Candidates strategy: reject all offers except the ones from their best possible stable partners, pretending that they prefer to not be matched rather than to the worse stable partners.

Manipulability in practice

Empirical studies [Roth and Peranson, 1999]
National Resident Matching Program: physicians in the USA.

"The set of stable matchings, and the opportunities for strategic manipulation, are surprisingly small."

A theoretical explanation

Theorem [Demange, Gale and Sotomayor, 1987]
Someone who belongs to exactly one stable pair cannot manipulate.

How can we explain this lack of manipulability ?

In the worst case: There exist setups s.t. every candidate, who normally would obtain her worst choice, has the opportunity to manipulate to obtain her best choice instead!

On average: When all preference lists are independent random uniform permutations, candidates have the opportunity to manipulate and, instead of getting a university of rank about n/log n, obtain a university of rank about log n

Consider "beyond worst case" settings

Correlated preferences and manipulability

Notions of (approximate) non-manipulability

1. If xxxx then everyone has a unique stable partner

2. If xxxx then almost everyone has a unique stable partner

3. If xxxx then for almost everyone, the stable partners are ranked roughly the same in their preference list

Goal: define xxxx so that
- we can prove the conclusion
- assumption xxxx is not unreasonable in practice

1.The case of common rankings

Fact: if all universities have the same preference list then the stable matching is unique (hence everyone has a unique stable partner, hence non manipulable)

Proof:

A:	1	2	3	4	5	6	none
B:	1	2	3	4	5	6	none
C:	1	2	3	4	5	6	none
D:	1	2	3	4	5	6	none
E:	1	2	3	4	5	6	none
F:	1	2	3	4	5	6	none

Candidate 1 must be matched to her favorite university x, o.w.:
if 1 is matched to y, then
1 prefers x to y
x prefers 1 to anyone,
so (1,x) is a blocking pair.

Iterating determines the stable matching

2. Approximate non-manipulability

Ashlagi Kanoria Leshno 2015: if there are N candidates and N-1 universities, and all preference lists are independent random uniform permutations, then the expected total number of stable partners is N+o(N).

Then almost everyone has a unique stable partner

3. Popularities

Fact: if all universities have the same preference list then the stable matching is unique

Intuition: rankings of students by universities are similar

Is the above fact robust when the preference lists are similar but not quite identical?

Roth Peranson 1999: "One factor that strongly influences the set of stable matchings is the correlation of preferences among programs and among applicants. When preferences are highly correlated (i.e., when similar programs tend to agree which are the most desirable applicants, and applicants tend to agree which are the most desirable programs), the set of stable matchings is small."

approximate non-manipulability:

xxxx = universities have highly similar preference lists, and candidates have highly similar preference lists

How do we model this family of inputs?

Popularity preferences [Immorlica Mahdian 2005]

Definition. Define the popularities p_A, p_B, p_C, p_D of A, B, C, D.

$$\mathbb{P}(D \succ A, B, C) = \frac{p_D}{p_A + p_B + p_C + p_D}$$

Sample using the popularities, without replacement.

$$\mathbb{P}(D \succ C \succ A \succ B) = \frac{p_D}{p_A + p_B + p_C + p_D} \cdot \frac{p_C}{p_A + p_B + p_C} \cdot \frac{p_A}{p_A + p_B}$$

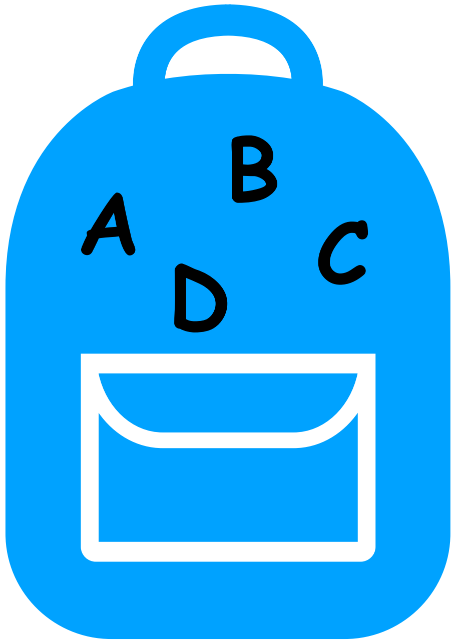

Preference list: D C A B

Geometric example

Assume the popularities of universities are

A:1,B:1/2,C:1/4,D:1/8,…

Then every candidate's preference list is similar to A B C D …

Q: What is the structure of stable matchings?

A: For each candidate, in expectation all of the stable partners are ranked within O(1) of one another in her preference list.

Remark: the assumptions are about the preference lists of candidates : the preference lists of universities are arbitrary!

Generalization of the geometric example

Definition: A distribution is regular if conditioning on {B>C>D>E>F} can only increase the probability that B is ranked ahead of other universities.

Theorem A. *Assume that each candidate independently draws her preference list from a regular distribution. The universities preference lists are arbitrary. Let x_k be an upper bound on the odds that university u_{i+k} is ranked before university u_i:*

$$\forall k \geq 1, \quad x_k = \max_{c,i} \left\{ \frac{\mathbb{P}[u_{i+k} >_c u_i]}{\mathbb{P}[u_i >_c u_{i+k}]} \right\}$$

Then for each candidate with at least one stable partner, in expectation all of her stable partners are ranked within $(1 + 2\exp(\sum_{k\geq 1} kx_k)) \sum_{k\geq 1} k^2 x_k$ of one another in her preference list.

Bound is interesting for distributions that produce preference lists that are all similar to a single "master" list, e.g. identical, geometric, from common+idiosyncratic utilities…

Theorem A answers the question of defining xxxx s.t.

"If xxxx then for almost everyone, the stable partners are ranked roughly the same"

Theorem A. *Assume that each candidate independently draws her preference list from a regular distribution. The universities preference lists are arbitrary. Let* x_k *be an upper bound on the odds that university* u_{i+k} *is ranked before university* u_i*:*

$$\forall k \geq 1, \quad x_k = \max_{c,i} \left\{ \frac{\mathbb{P}[u_{i+k} >_c u_i]}{\mathbb{P}[u_i >_c u_{i+k}]} \right\}$$

Then for each candidate with at least one stable partner, in expectation all of her stable partners are ranked within $(1 + 2\exp(\sum_{k\geq 1} kx_k)) \sum_{k\geq 1} k^2 x_k$ *of one another in her preference list.*

4. A more structured example: geometric-uniform

Assume the popularities of universities are

A:1,B:1/2,C:1/4,D:1/8,…

Then every candidate's preference list is similar to A B C D …

Additionally, assume that the universities preference lists are i.i.d. random uniform permutations of the candidates.

Q: What is the structure of stable matchings?

A: Almost every person has a unique stable partner!

Generalization of the geometric-uniform example

Theorem B. *Assume that each candidate independently draws her preference list from a regular distribution. The universities preference lists are arbitrary. Let x_k be an upper bound on the odds that university u_{i+k} is ranked before university u_i:*

$$\forall k \geq 1, \quad x_k = \max_{c,i} \left\{ \frac{\mathbb{P}[u_{i+k} >_c u_i]}{\mathbb{P}[u_i >_c u_{i+k}]} \right\}$$

Further assume that $u_k = \exp(-\Omega(k))$, and that universities have uniformly random preferences. Then, in expectation the fraction of persons who have multiple stable partners converges to 0.

Theorem B answers the question of defining xxxx s.t.

"If xxxx then almost everyone has a unique stable partner"

Technical part 2
Proof of the geometric result

Universities/Men are matched to Candidates/Women

Geometric example

Assume the men's popularities are

A:1, B:1/2, C:1/4, D:1/8, …

Then every woman's preference list is similar to A B C D …

Q: What is the structure of stable matchings?

A: For each woman, in expectation all of the stable partners are ranked within O(1) of one another in her preference list.

Remark: the men's preference lists are arbitrary!

Analysis of the geometric example: Separators

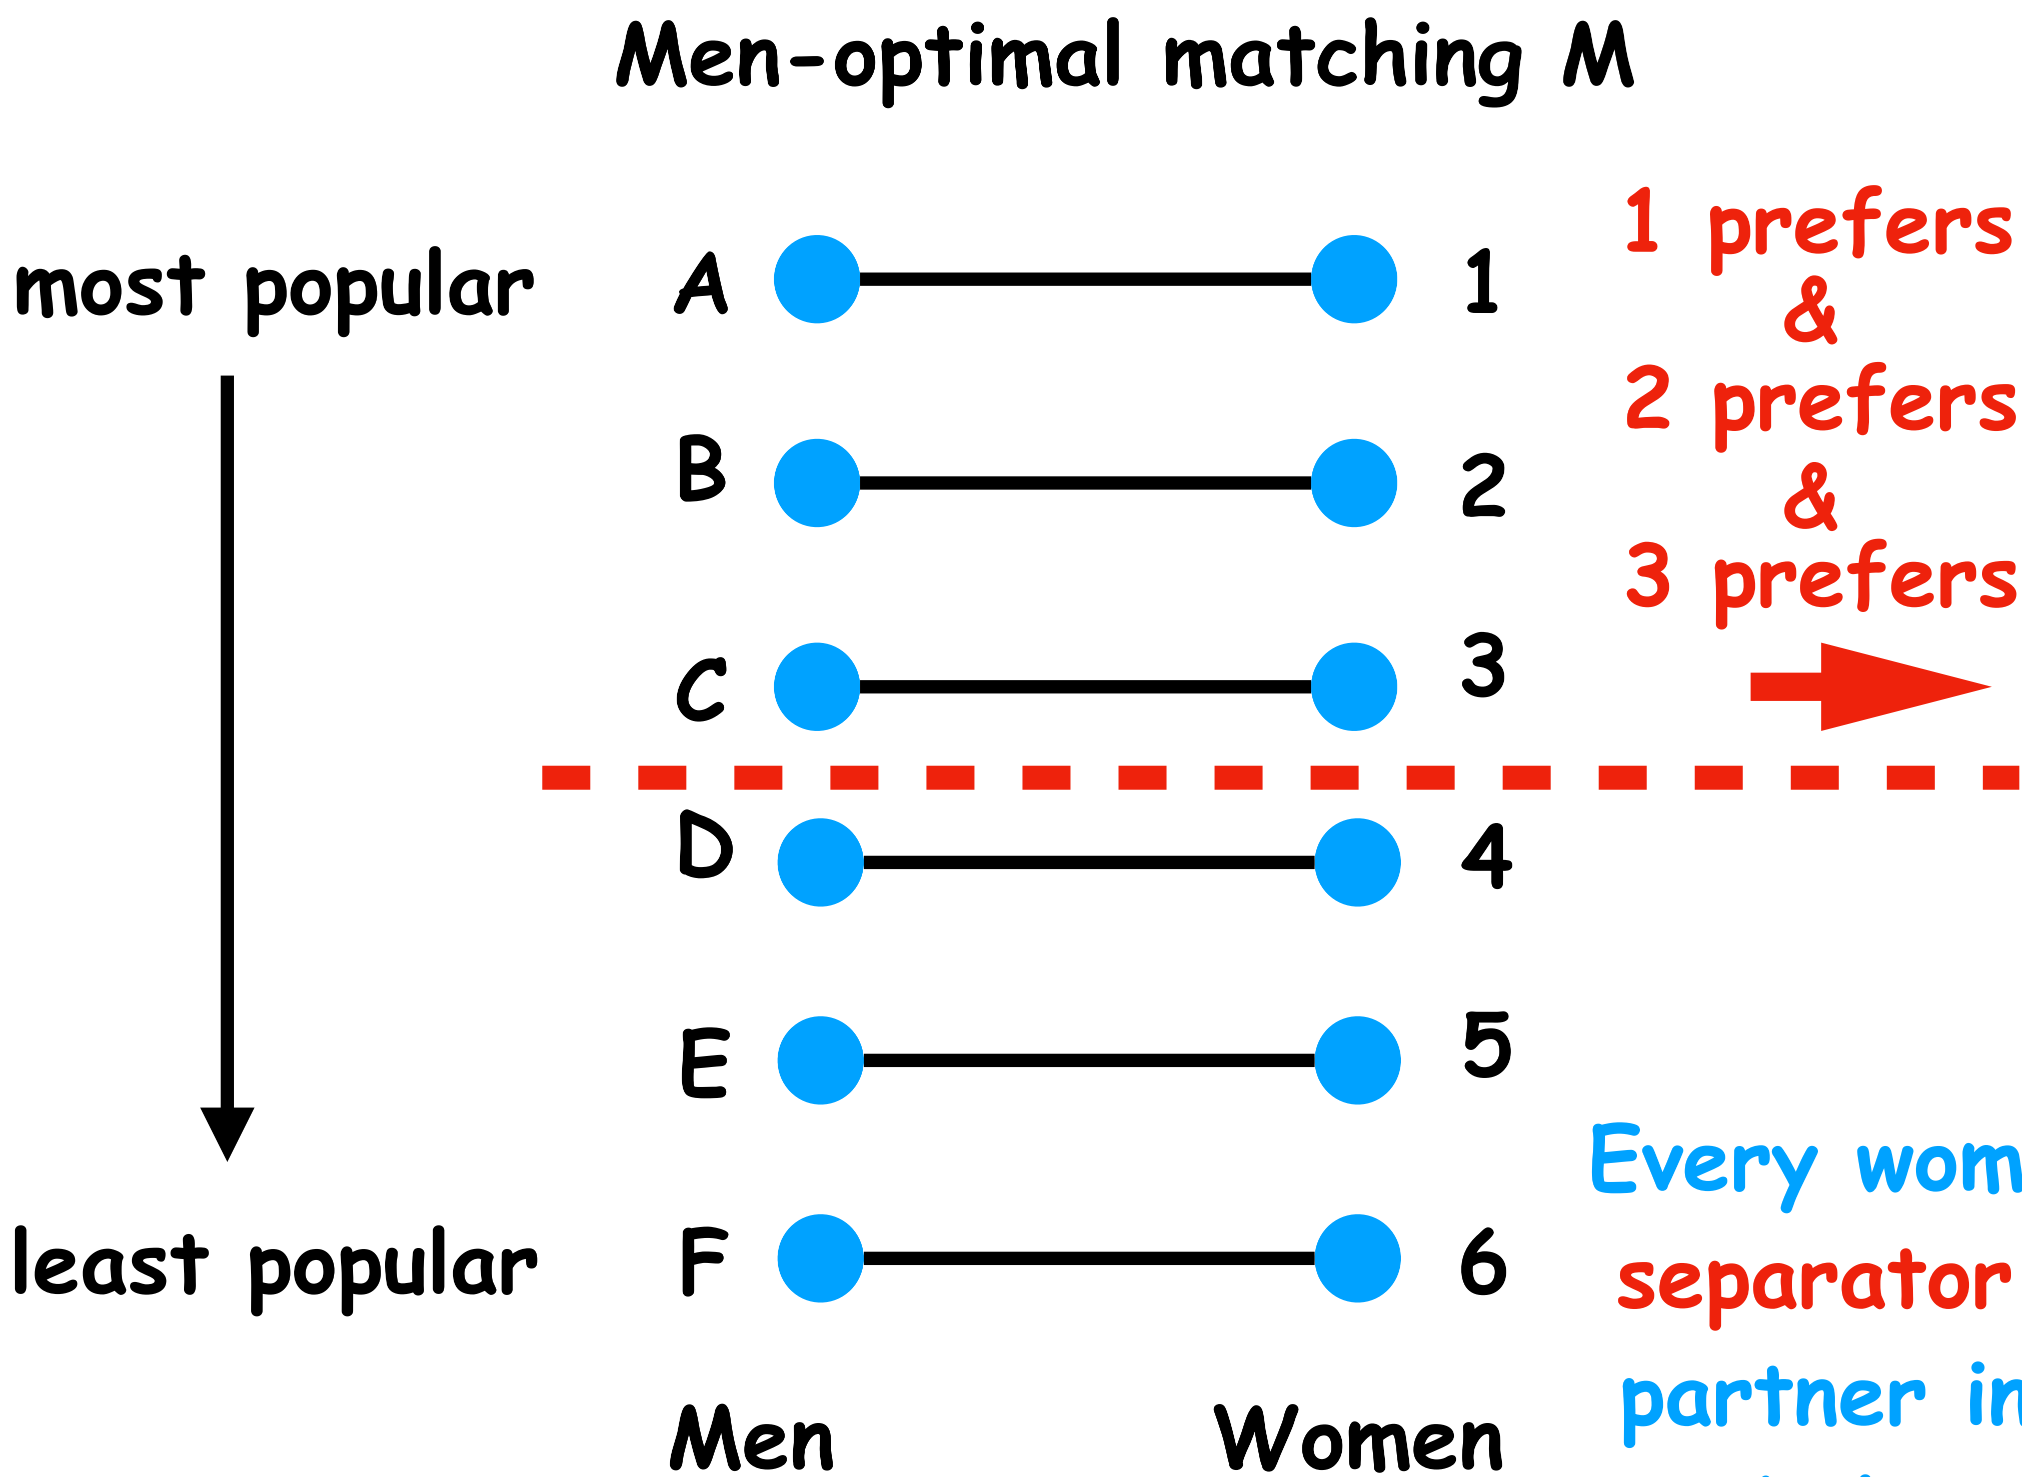

1 prefers A to D,E,F,...
&
2 prefers B to D,E,F,...
&
3 prefers C to D,E,F,...

separator

Every woman above the separator prefers her partner in M to every man below the separator

Claim: in every stable matching, {A,B,C} are matched to {1,2,3}

Men-optimal matching M

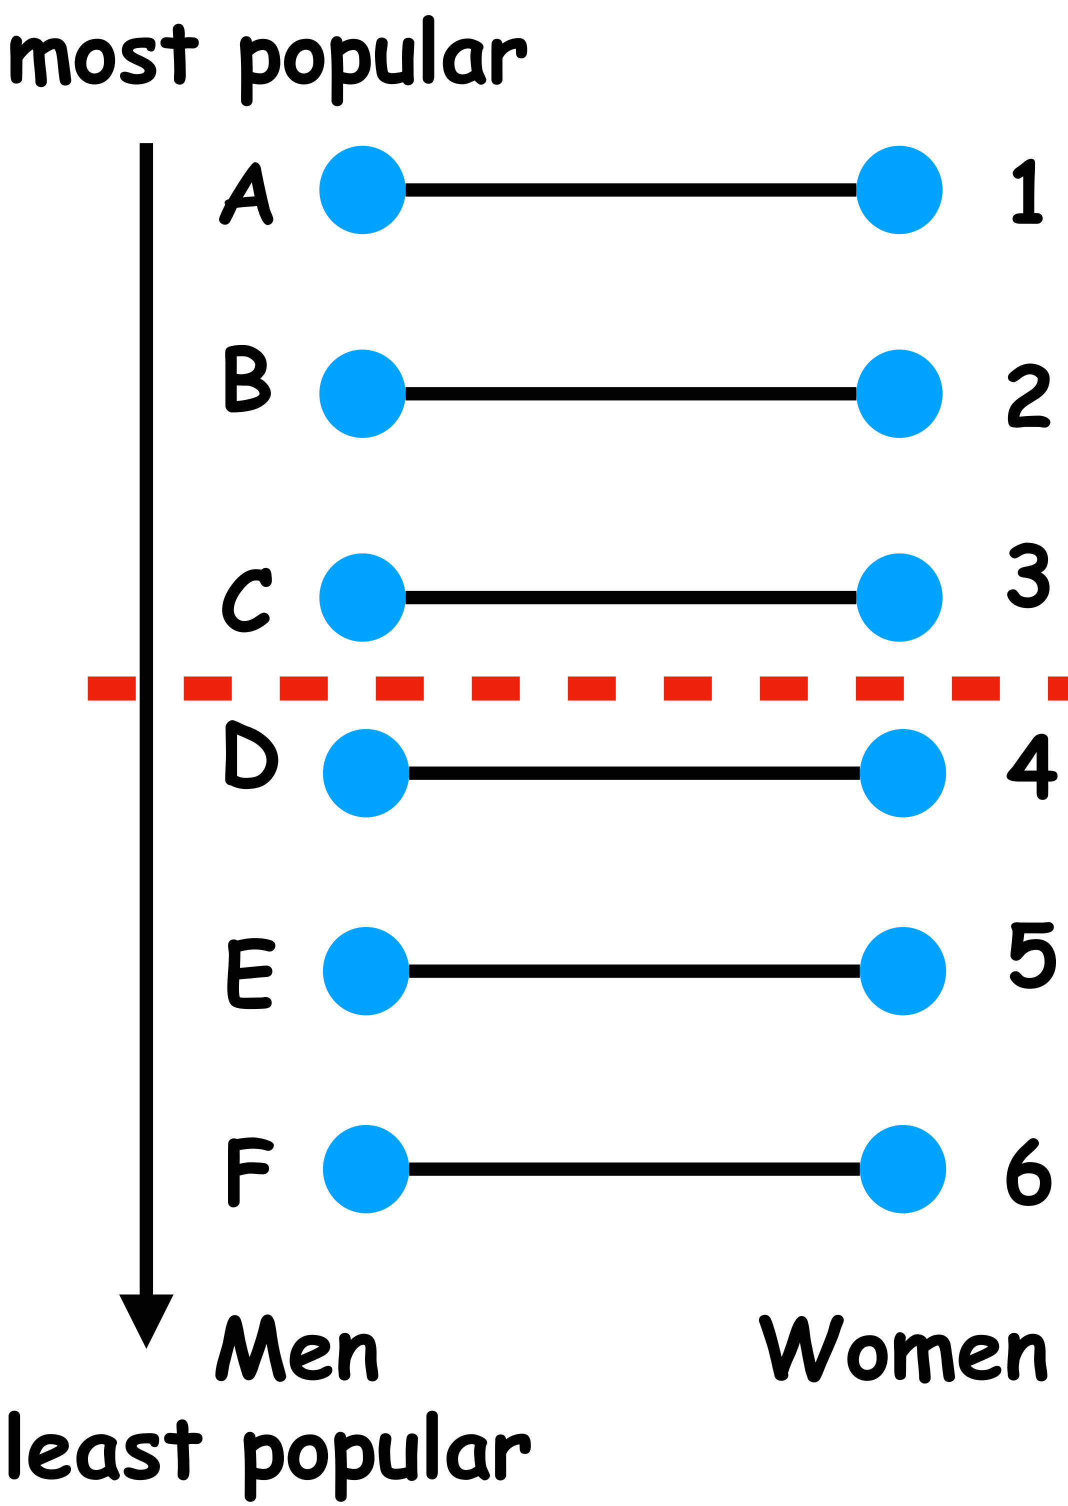

Pf: by contradiction.

Assume some stable matching M'
matches 2 to a man outside {A,B,C}.

Separator => 2 prefers B to her match in M'.

M man-optimal => B prefers 2 to his match in M'.

(B,2) is a blocking pair, contradiction.

Analysis of the geometric example: Blocks

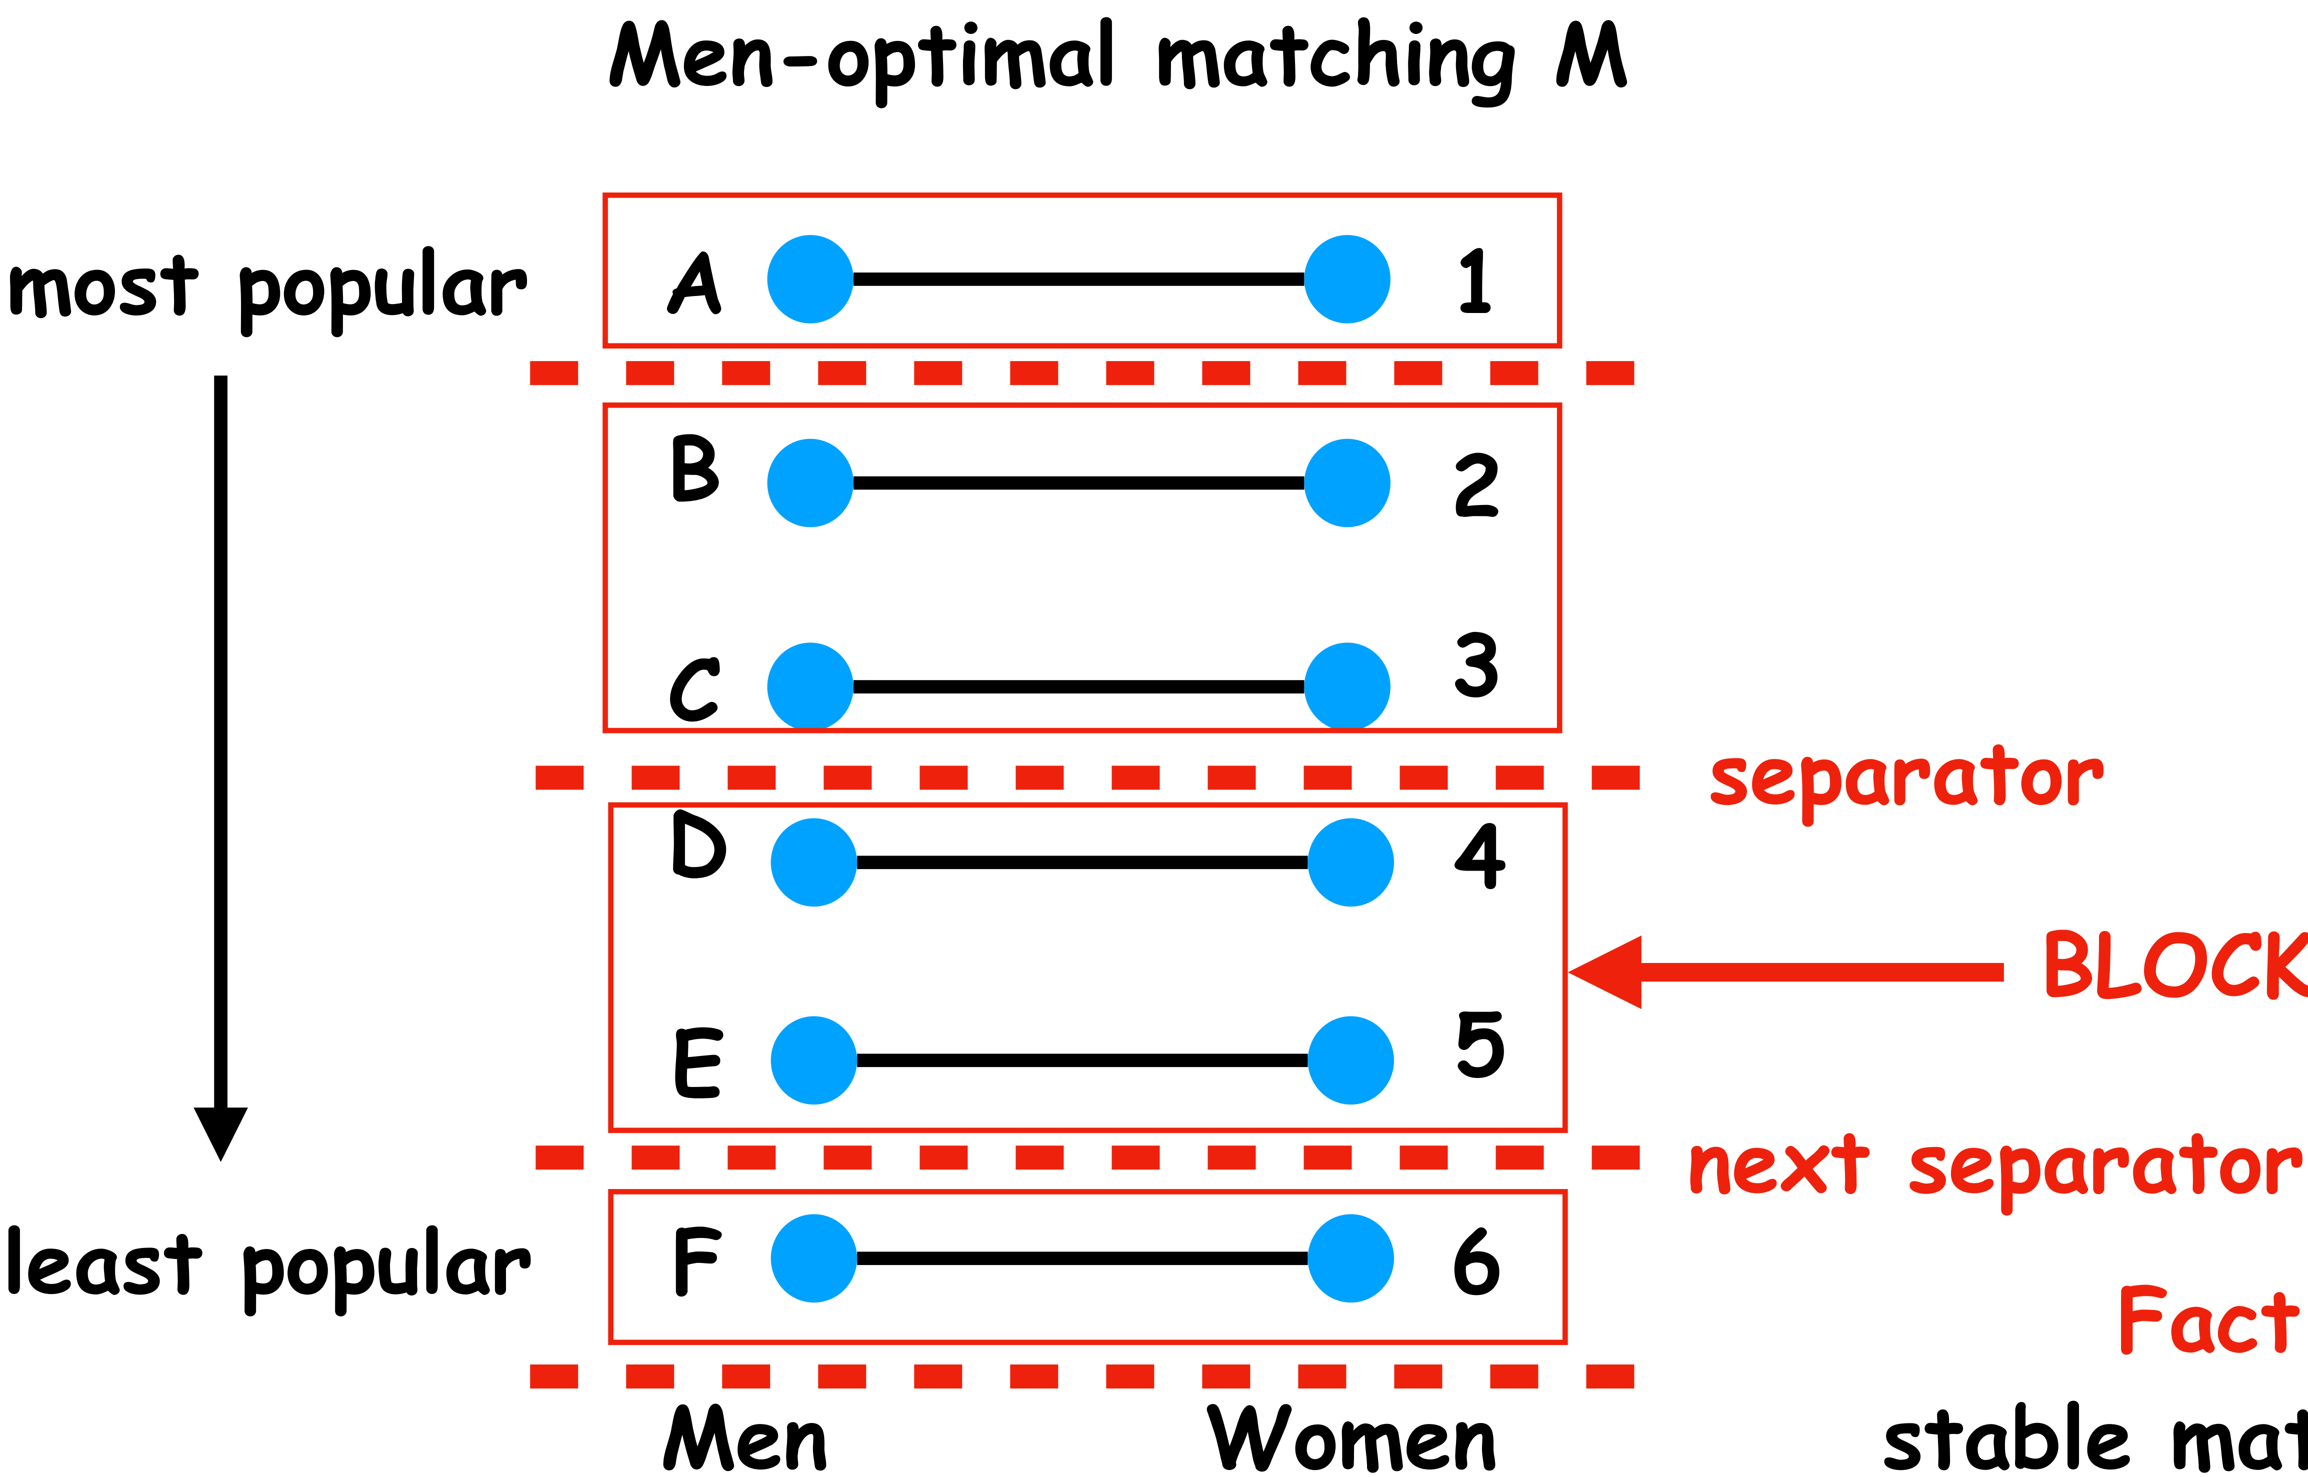

Fact:
stable matchings
stay within blocks

Analysis of the geometric example: structural lemma

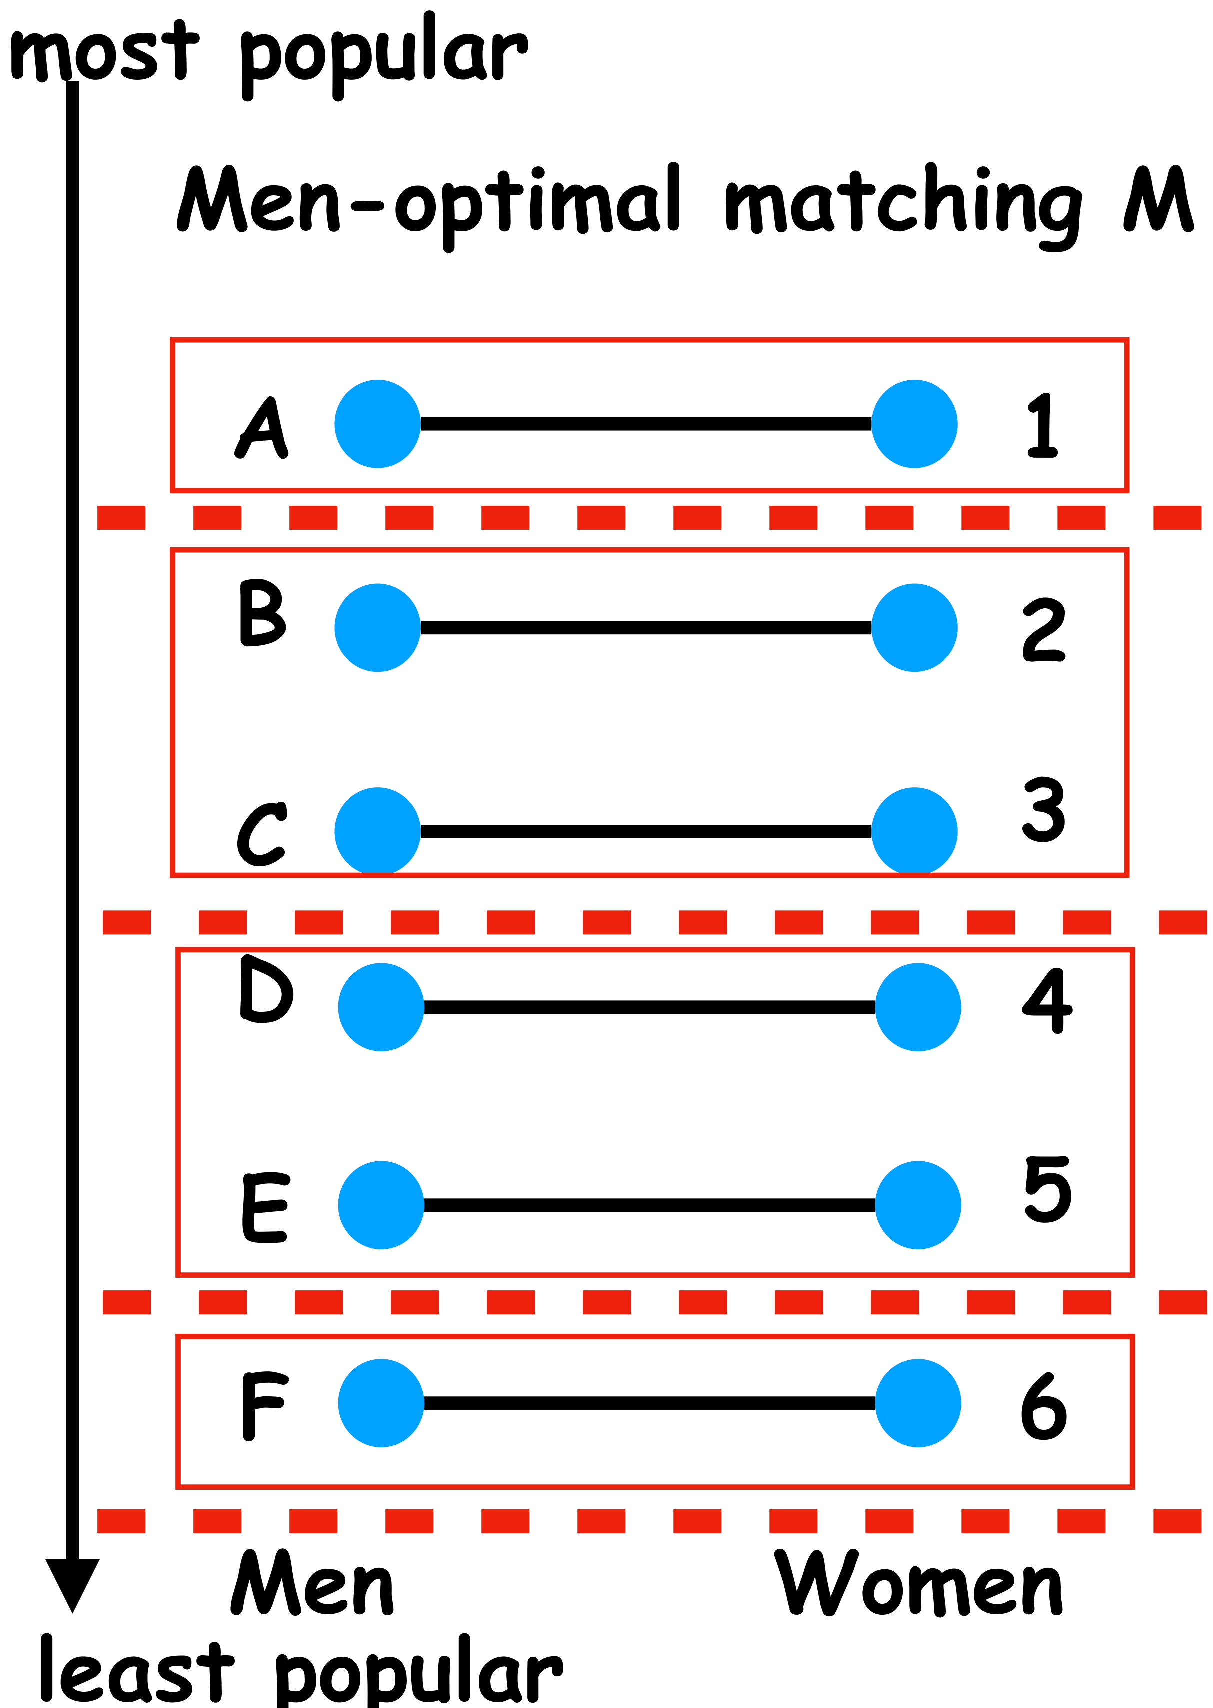

Focus on woman 2

Her man-optimal partner: B

Her block: {B,C}

Her stable partners:

subset of {B,C}

Her preference list:

B A F C D E

Difference in rank between her best and worst stable partners : at most

(size of her block) + #(men inserted in-between in her list)

Rank diff between best and worst stable partners of woman w

at most : (size of her block) + #(men inserted in-between in her list)

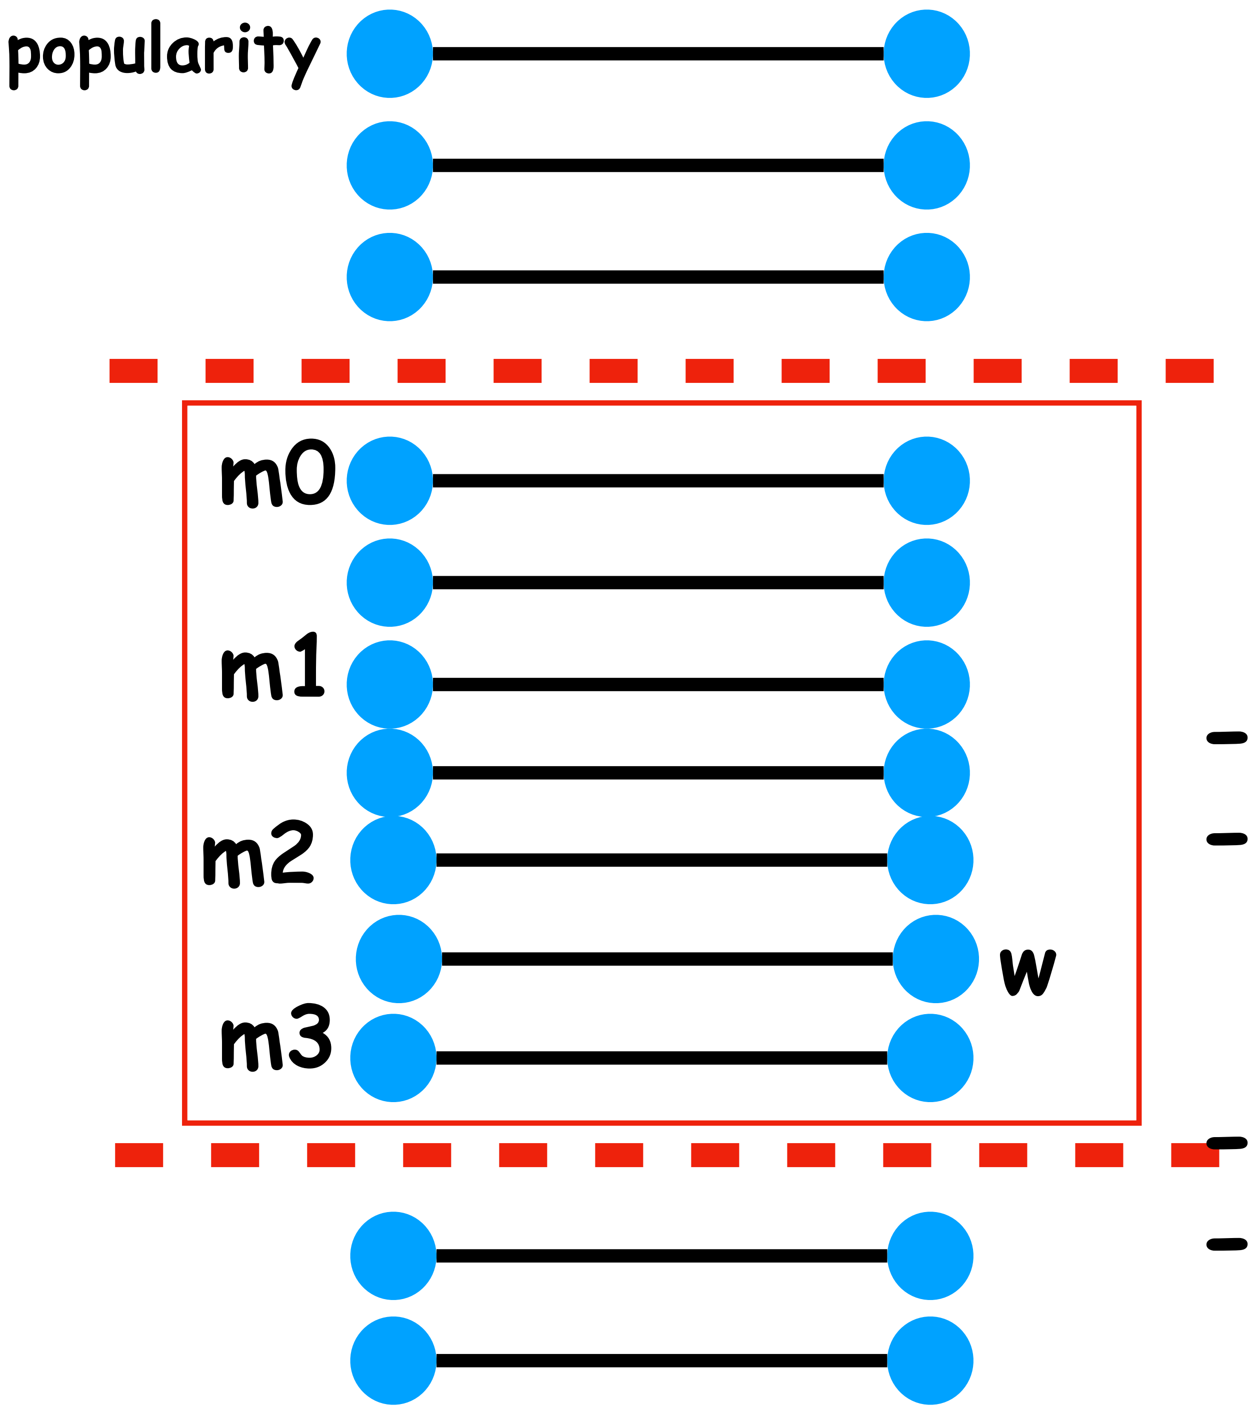

Pref list of w:

… m1 … m0 … m3 … m2 …

All of w's block is between m1 and m2

Since the list is drawn using geometric popularities 1,1/2,1/4,…, most likely:

- m0 is not far behind m1
- few men in previous blocks are ranked after m1

Similarly:

- m3 is not far before m2
- few men in later blocks are ranked before m2, so:

#(men inserted in-between in her list) = O(1)

Rank diff between best and worst stable partners of woman w

at most : (size of her block) + #(men inserted in-between in her list)

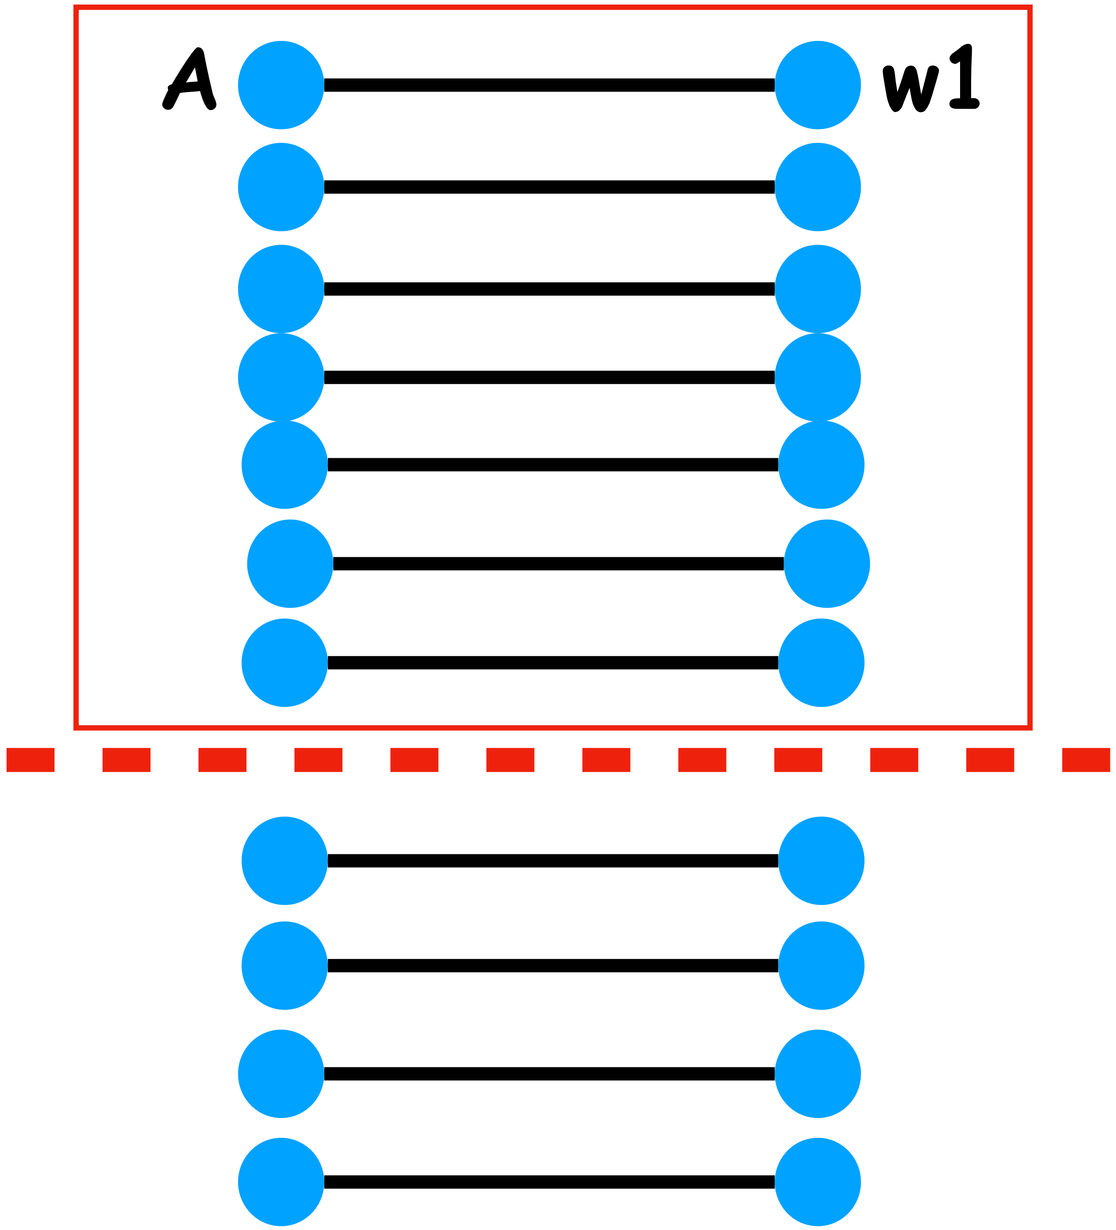

Analyzing block sizes:

- analyze size of first block
- iterate

Algorithm to construct first block

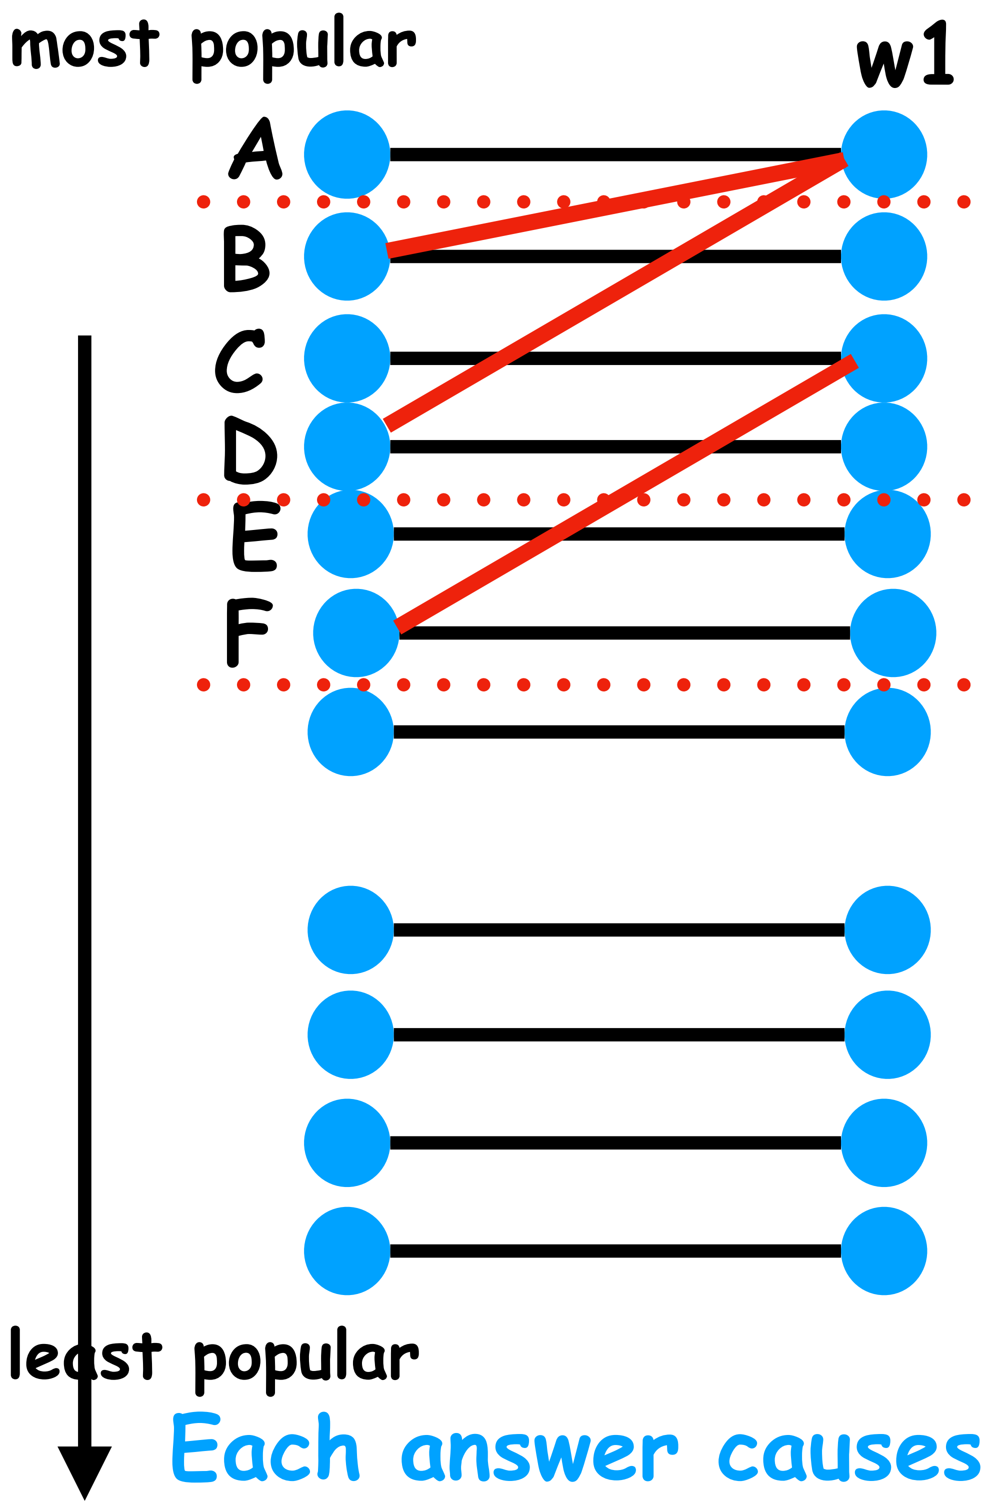

Q1: Who does w1 prefer to A?

A1: B and D.

Q2: Outside [A,D], who do [w1,w4] prefer to their partner?

A2: F.

Q3: Outside [A,F], who do [w1,w6] prefer to their partner?

A3: no one.

Then first block = [A,F].

Analysis:

Each answer causes a jump down the list of O(1) on average

Each question has constant of receiving the answer "none"

so there are O(1) iterations.

So block size=O(1)*O(1)=O(1).

Analysis of the geometric example, concluded

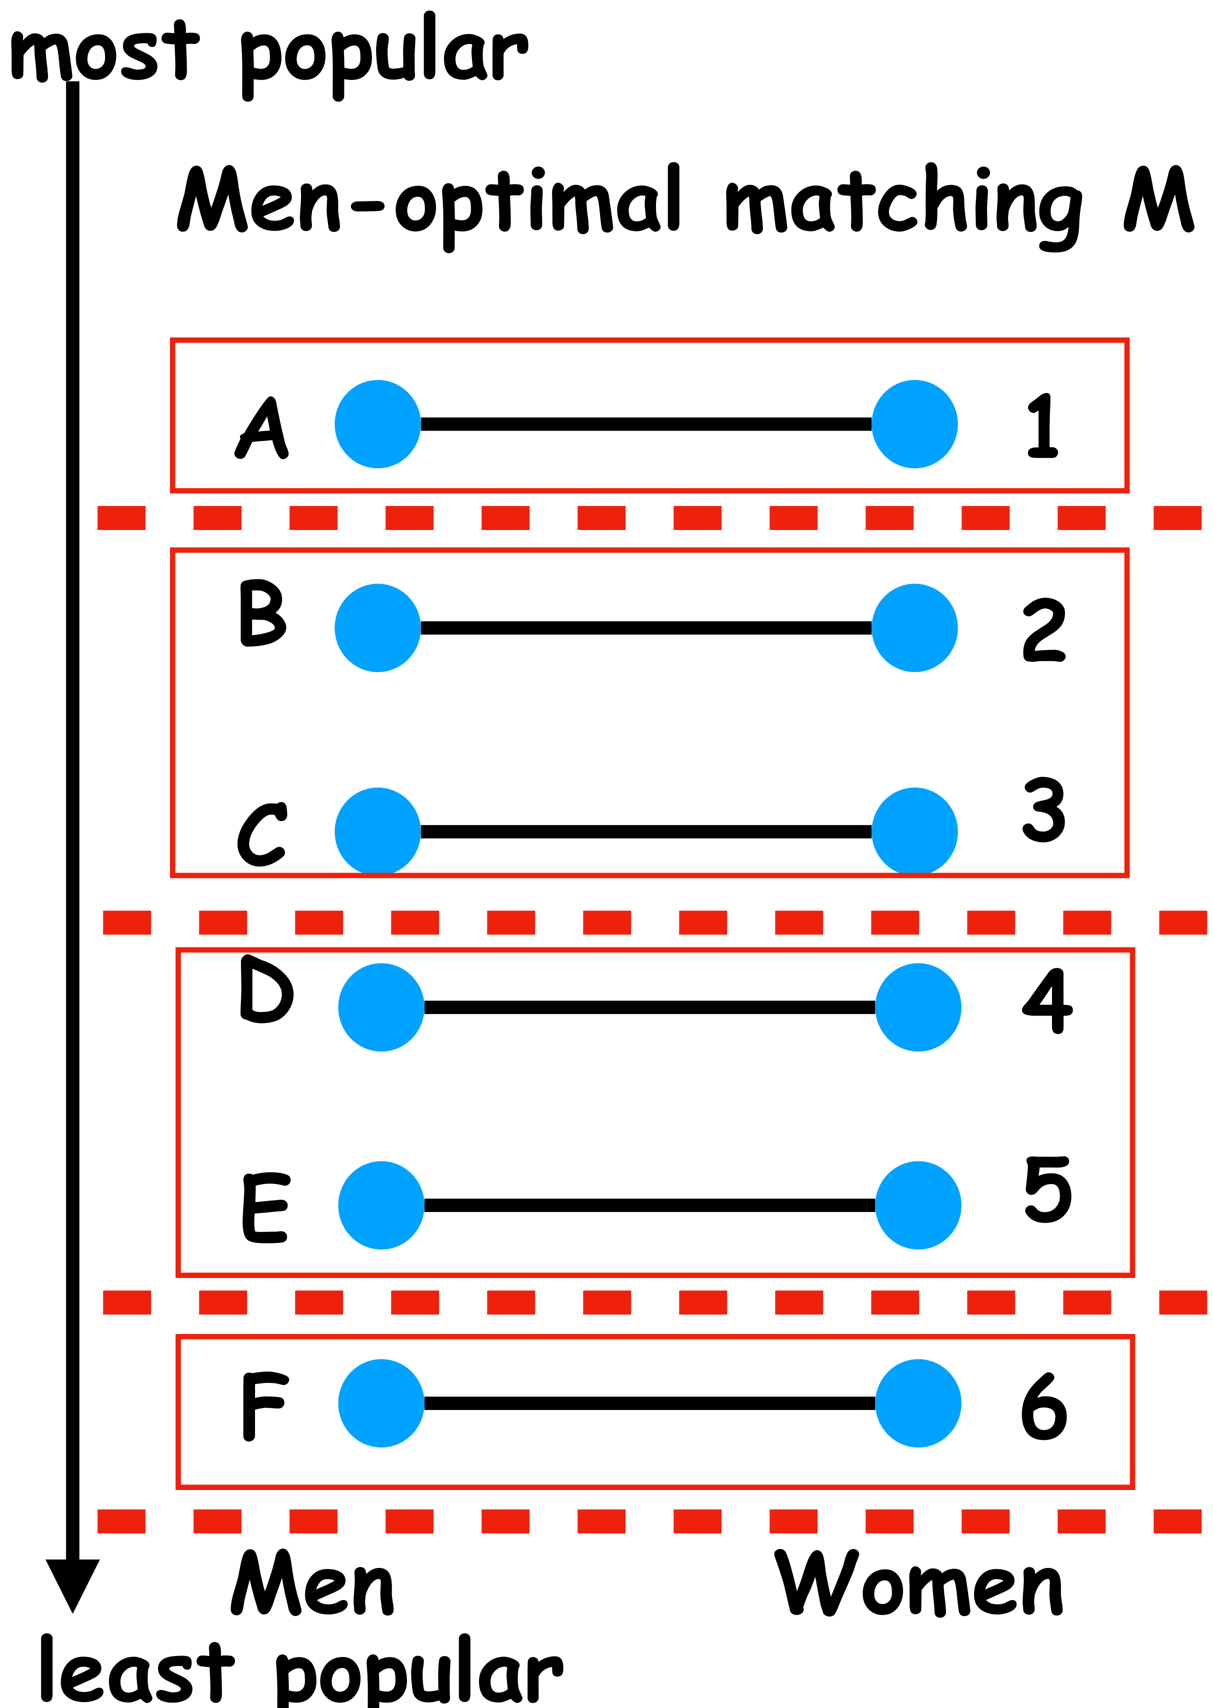

Focus on woman w

Difference in rank between her best and worst stable partners : at most

(size of her block) +
#(men inserted in-between in her list)
=
O(1)+O(1)=O(1).

Recap of what we have proved

Geometric example

Assume the men's popularities are

A:1, B:1/2, C:1/4, D:1/8, …

Then every woman's preference list is similar to A B C D …

Q: What is the structure of stable matchings?

A: For each woman, in expectation all of the stable partners are ranked within $O(1)$ of one another in her preference list.

(**Remark**: the men's preference lists are arbitrary!)

=> approximate non-manipulability

Conclusion

Methodology

1. Start from a real-life or intuitive observation on a parameter of interest, manipulability
2. "the opportunities for strategie manipulation are surprisingly small"
3. Observe that the observation does not match the worst case

5. Observe that average case is an unreasonable assumption and/or leads to unsatisfactory conclusions

7. Study a "beyond worst case" model
8. Prove results on non-manipulability

Technical part 2
bound on the individual number of stable pairs

Theorem D. *Assume that:*

- *universities and candidates have popularity preferences; and*
- *the maximal ratio between the popularities of two candidates for the same university is polynomial in N; and*
- *the candidates do not disagree much on the relative popularities of universities:*

$$\max_{c_0,c_1}\left(\left(\max_u \frac{\mathcal{D}_{c_0}(u)}{\mathcal{D}_{c_1}(u)}\right)\cdot\left(\max_{u'} \frac{\mathcal{D}_{c_1}(u')}{\mathcal{D}_{c_0}(u')}\right)\right)$$

is polylog in N.

Then the expected number of stable matches of any candidate c is at most polylog in N.

Candidate 1 declines all proposals

Preferences are sampled online

Stochastic process

Start from the university proposing stable matching.
Candidate c **declines** all proposals.
Preferences are sampled **online.**

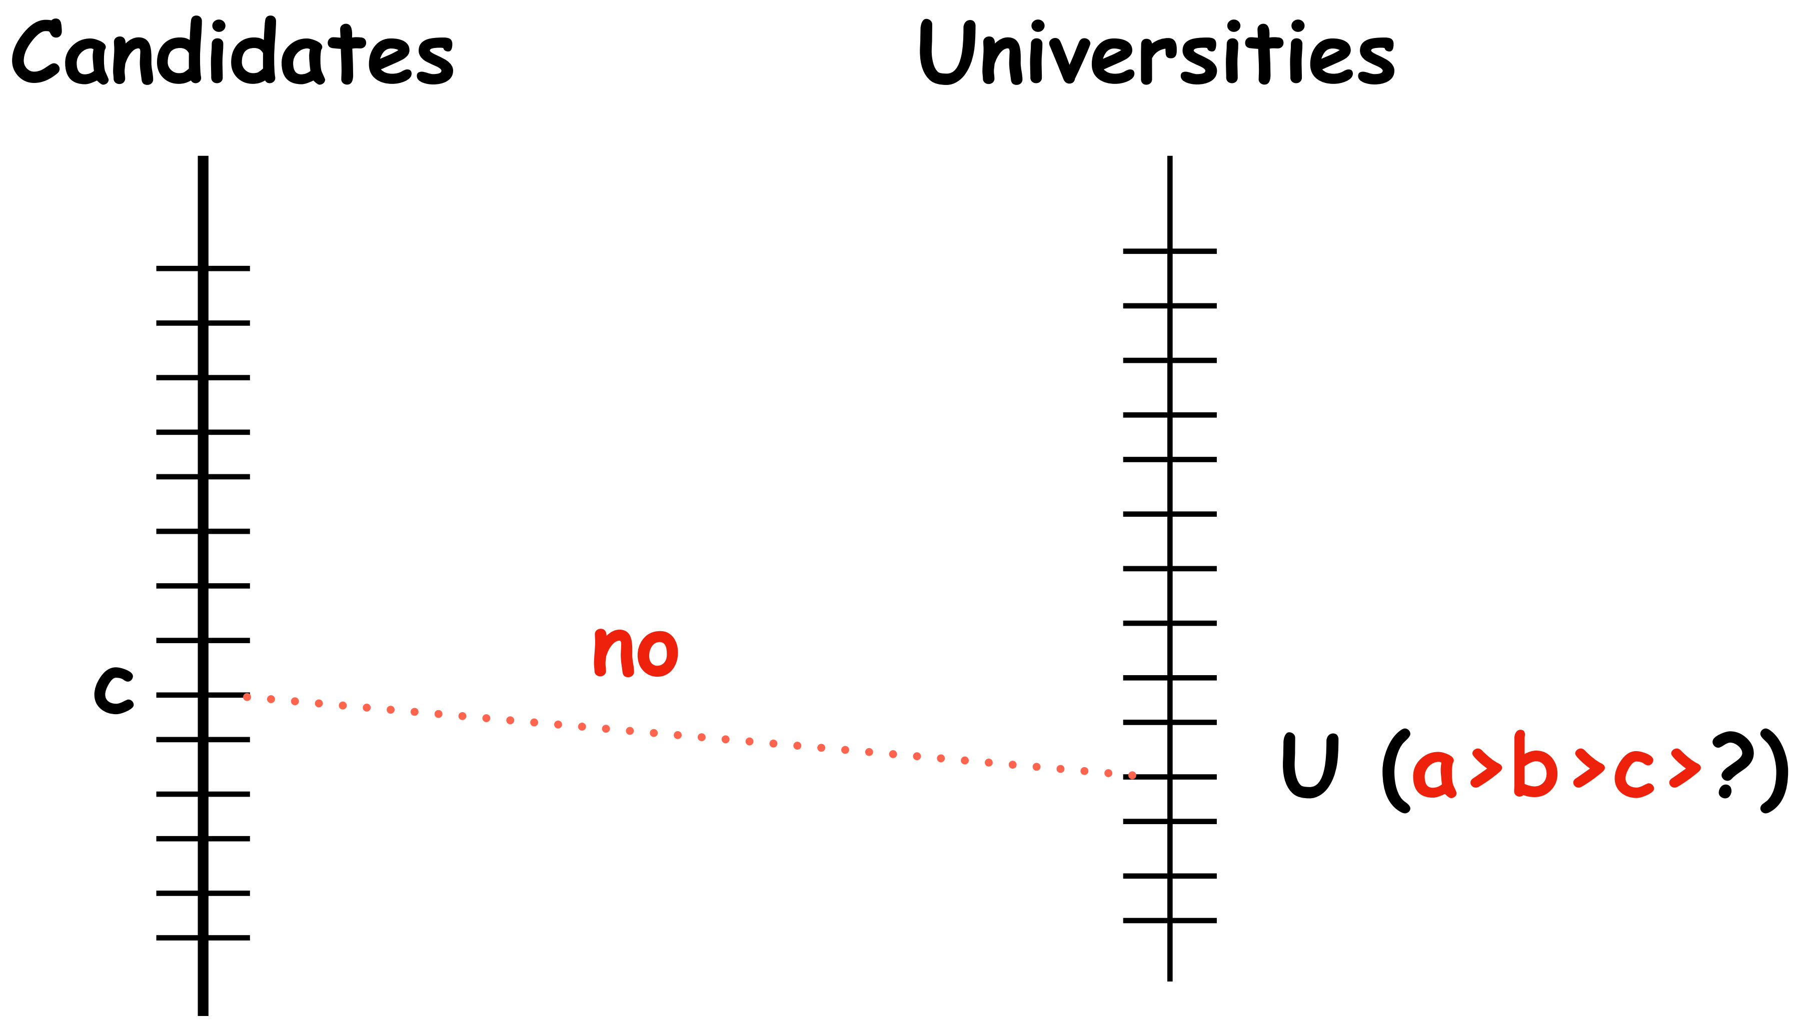

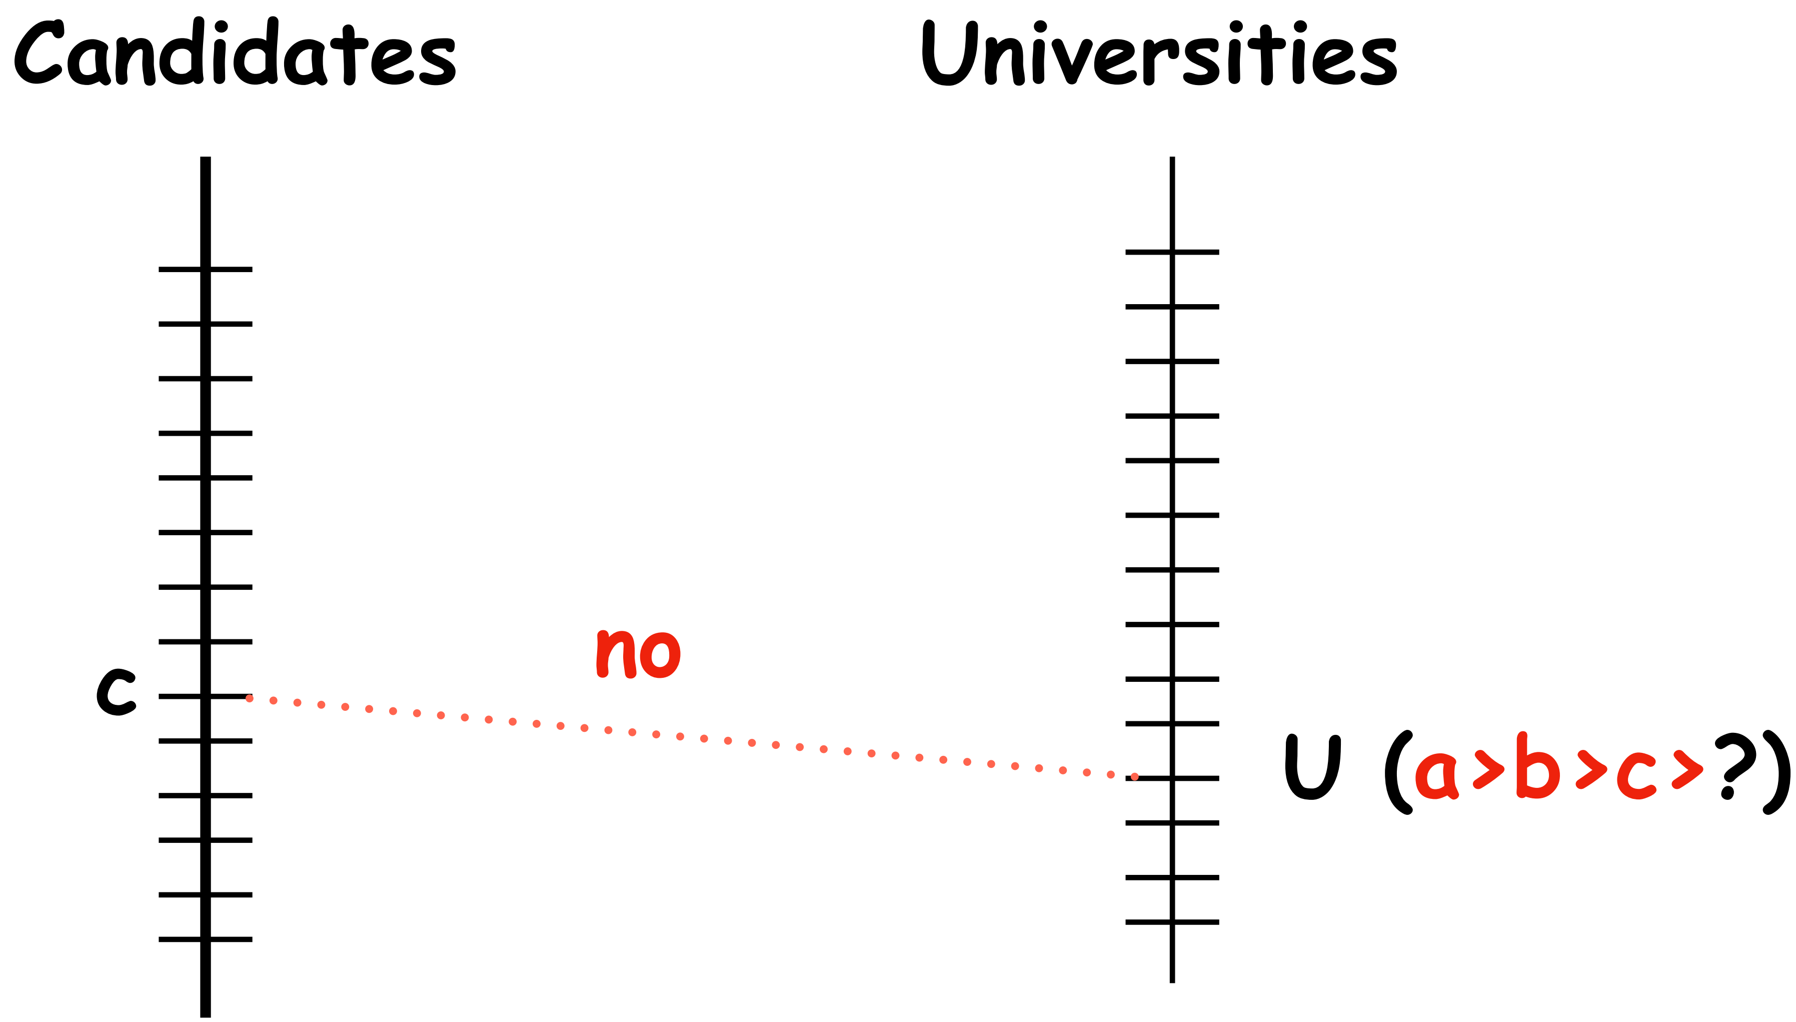

The next candidate d called by U is chosen with probability proportional to its popularity

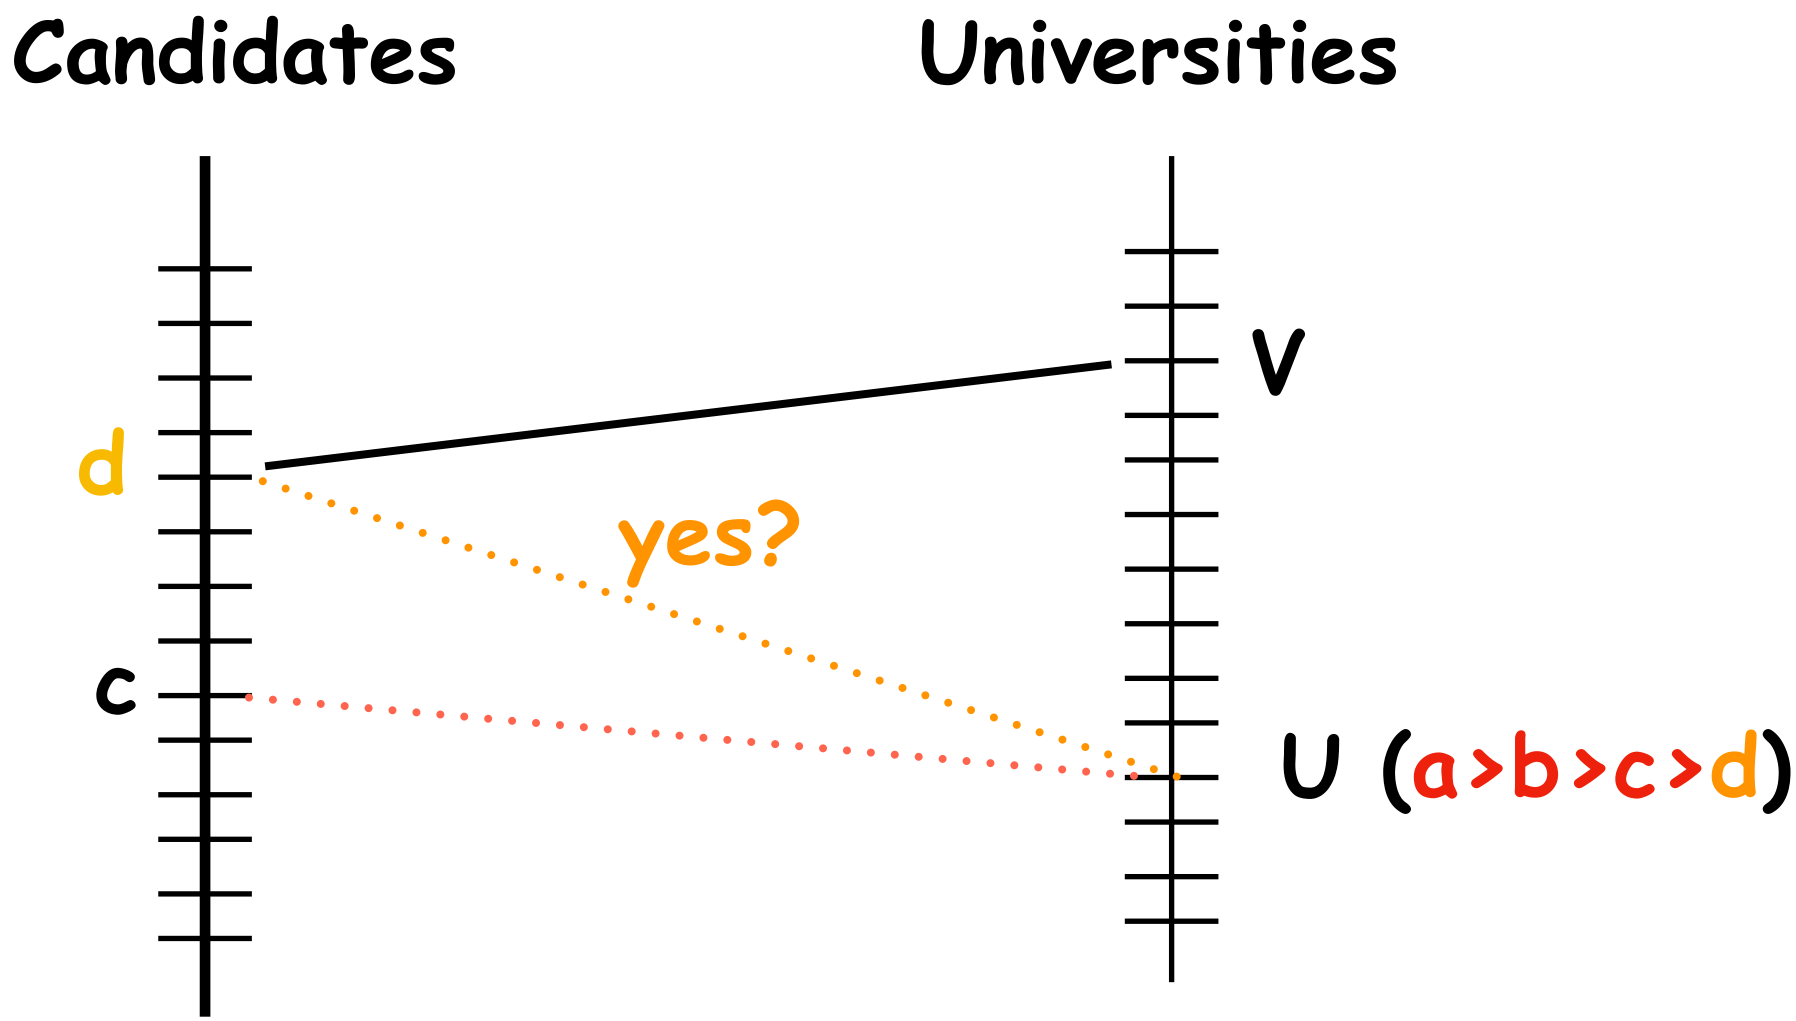

Probability that d says "yes" to U?

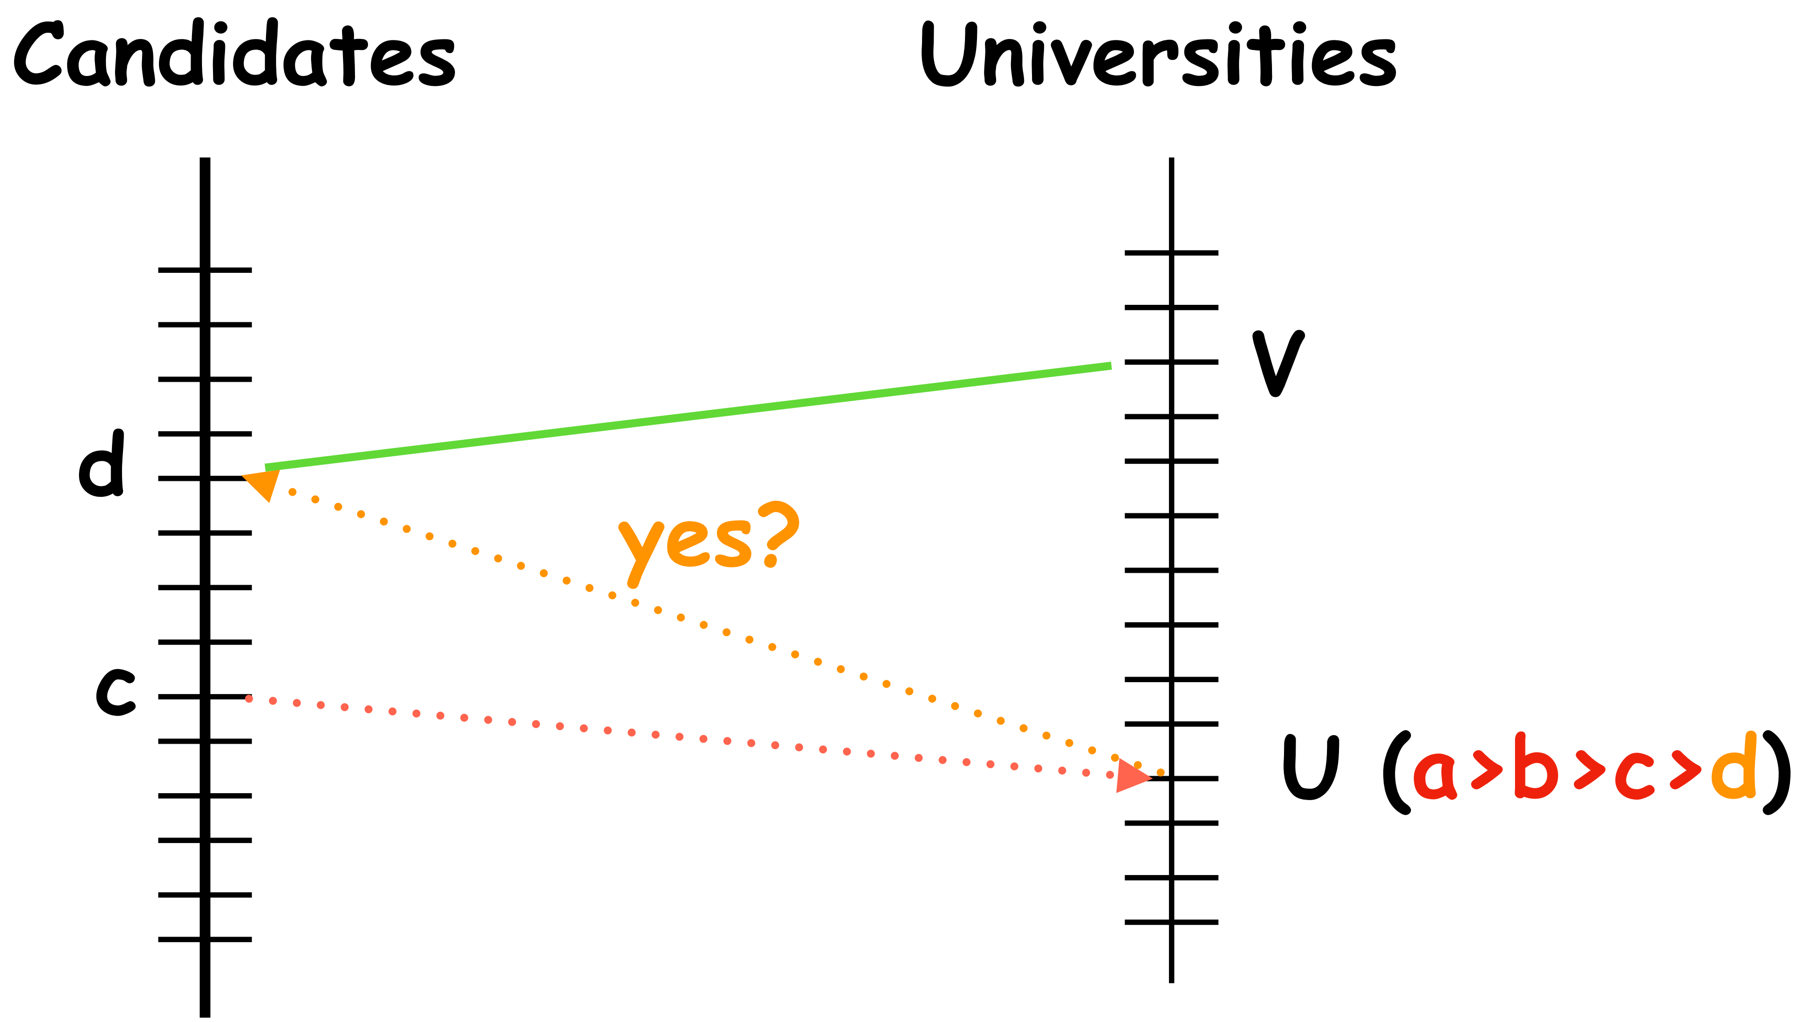

Probability that d says "yes" to U?

d should prefer U of popularity p(U) to all other proposals she got so far

Denote p(tot) the sum of popularities of universities having proposed to d

The answer of d is yes with probability

$$p(U) / (p(U) + p(tot))$$

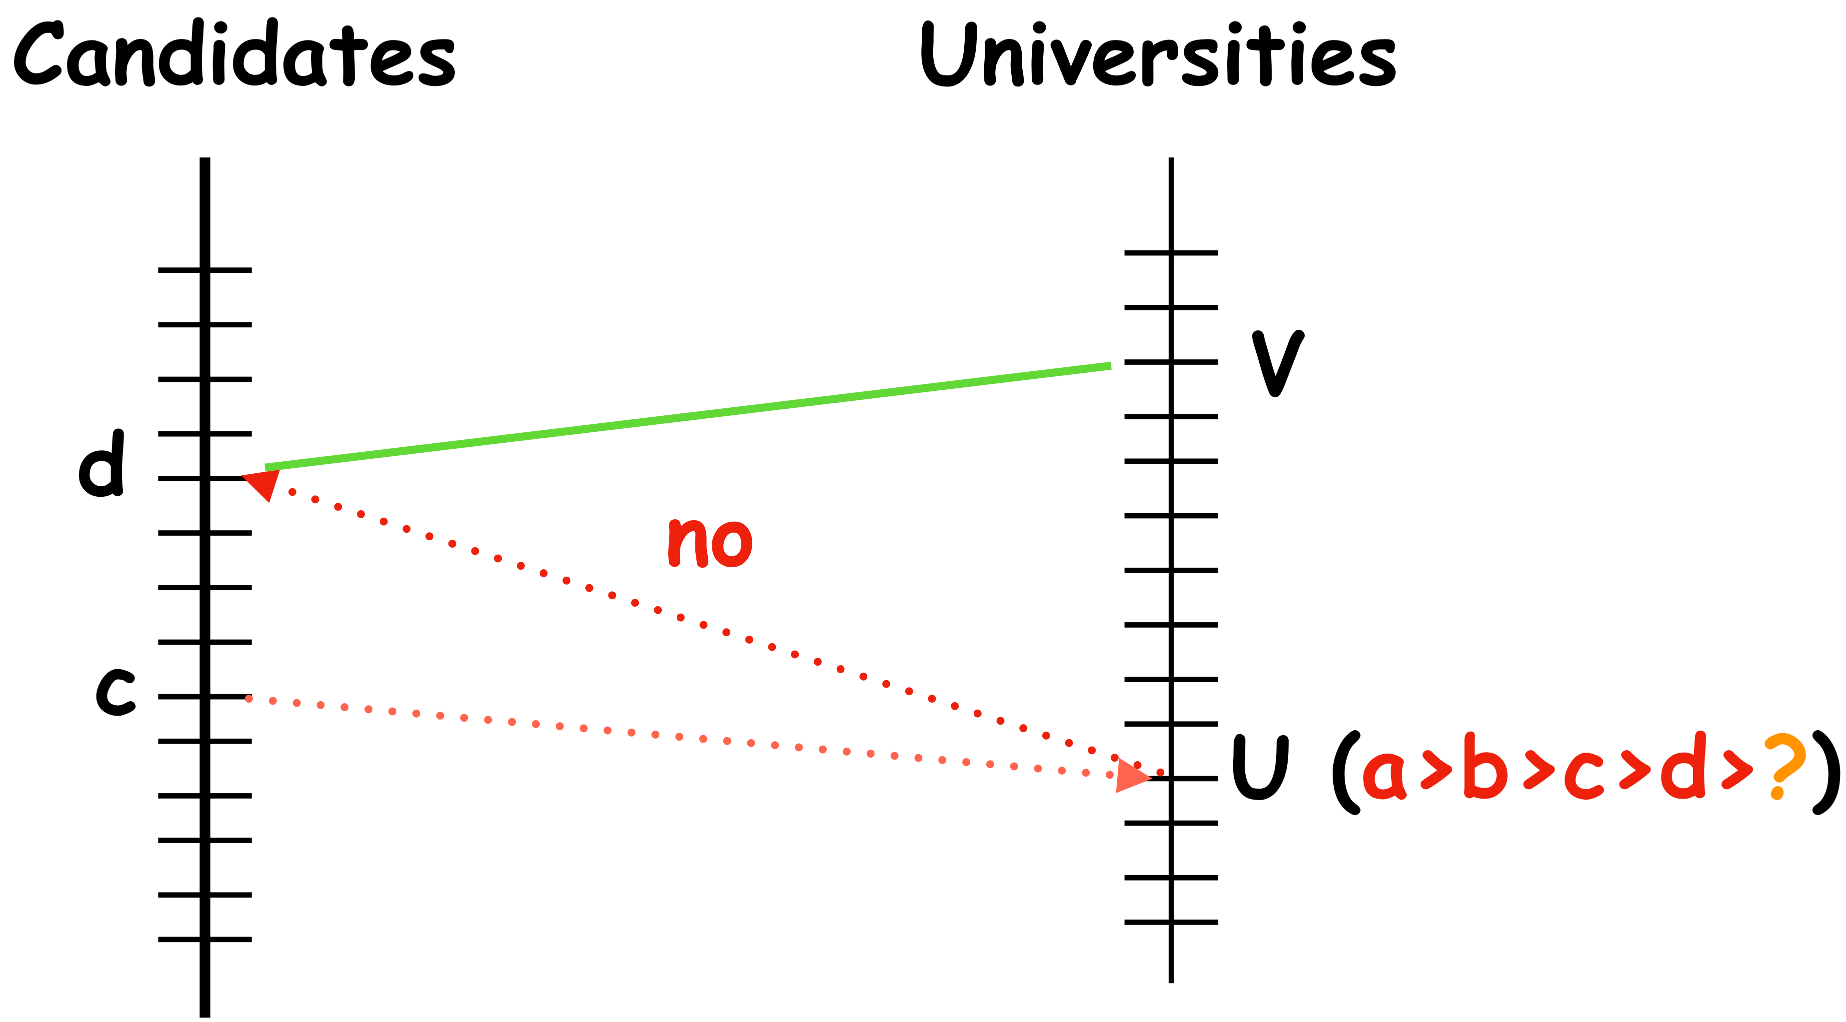

Probability that d says "yes" to U?

d should prefer U of popularity p(U) to all other proposals she got so far

Denote p(tot) the sum of popularities of universities having proposed to d

The answer of d is yes with probability

$$p(U) / (p(U) + p(tot))$$

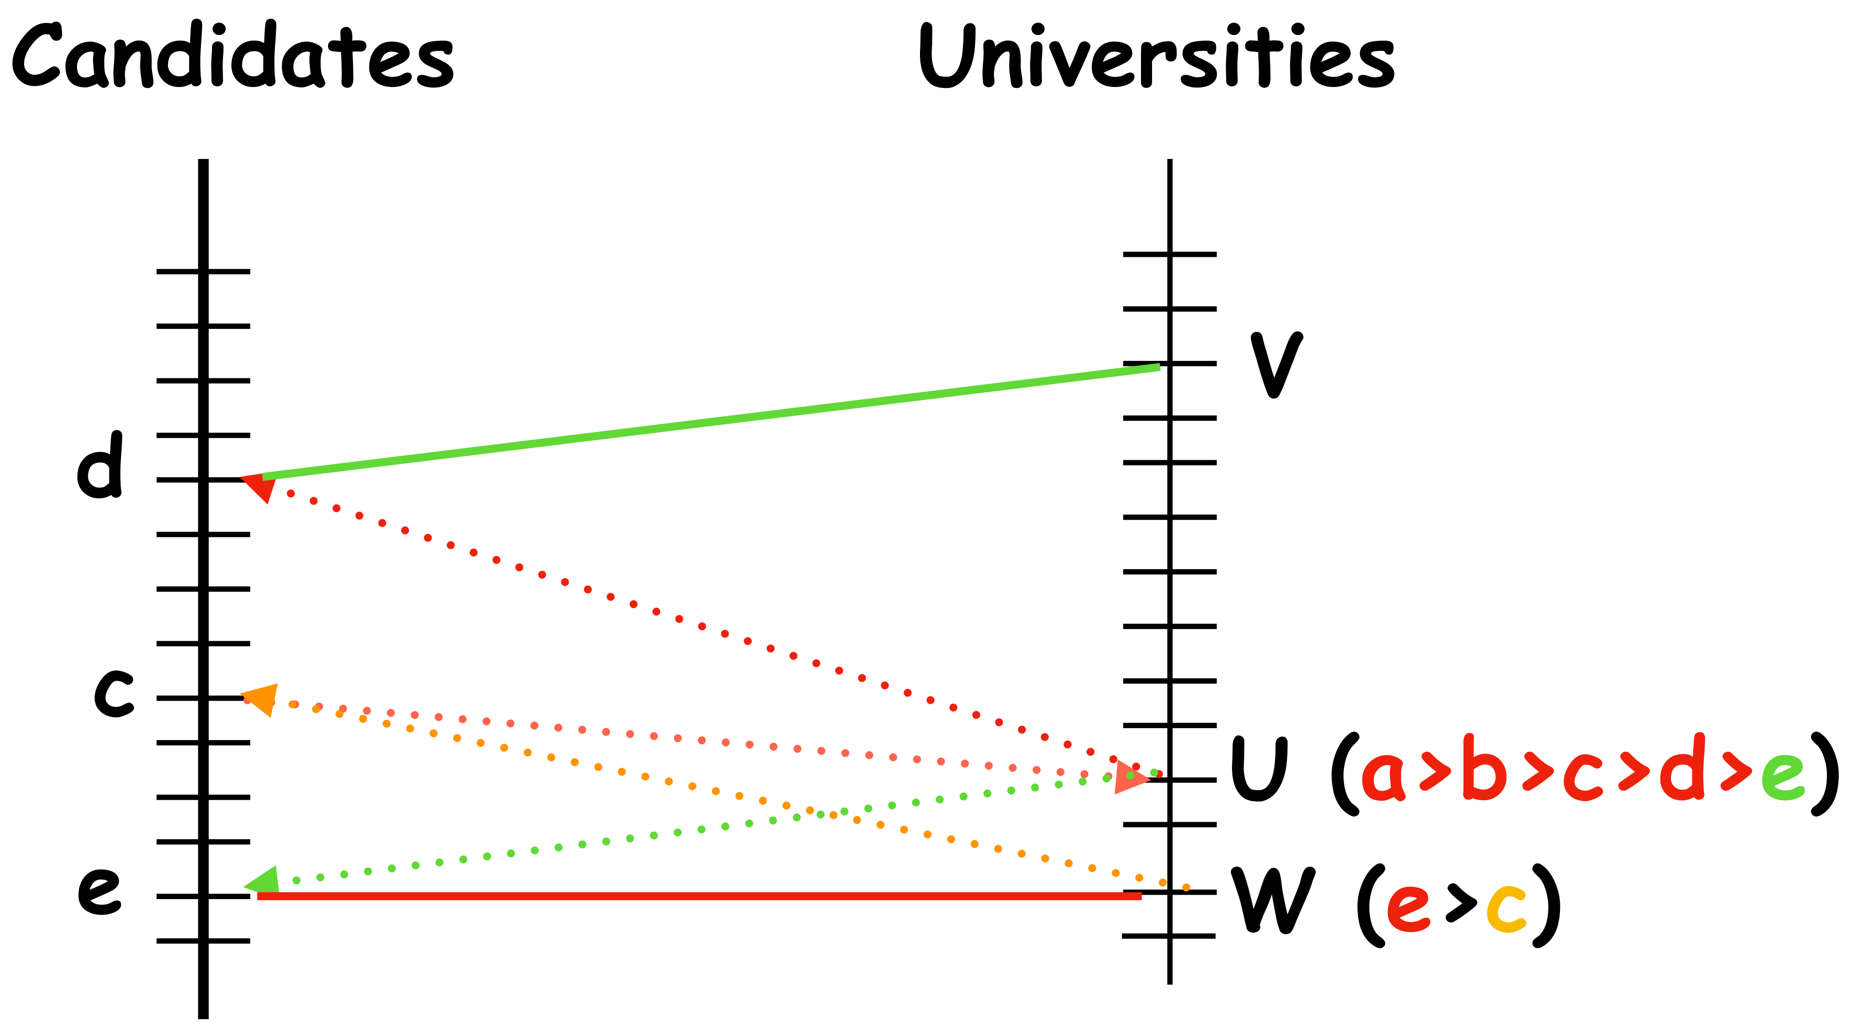

W is a new stable match for c

Possible stable matches of c?

Candidates almost agree on popularities of universities:
No large jumps on the right side.

For each university, candidates have similar popularities, up to poly(N)
Homogenous choice of the right-to-left moves

The random walk does not reach top universities
-> bound on the number of "yes" by c.