Strategy-Proofness in the Stable Matching Problem with Couples

Andrew Perrault, Joanna Drummond, and Fahiem Bacchus

Computer Science
UNIVERSITY OF TORONTO

Stable Matching Problem (SMP)

- Two-sided matching problem

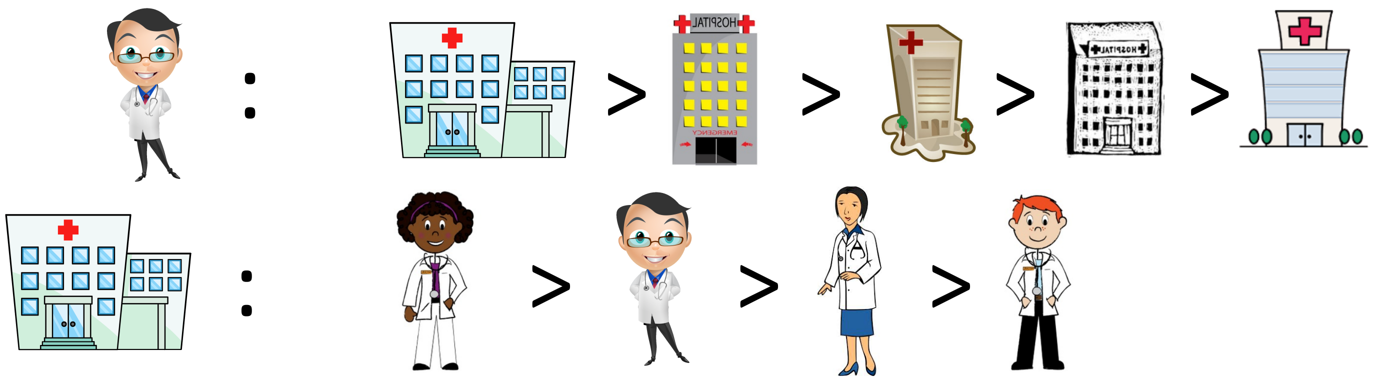

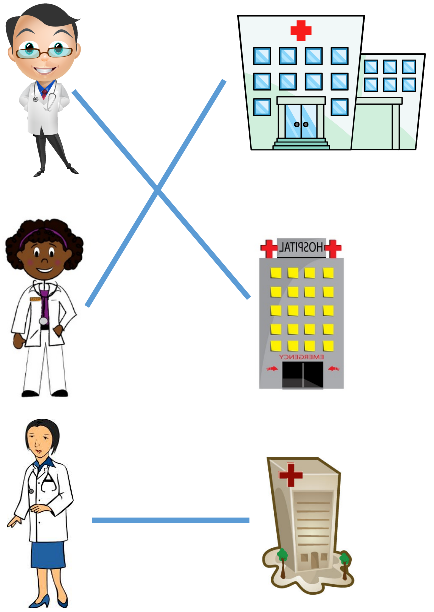

- *Stable matching:* no resident-hospital pair prefers each other to their current matching
- Polynomial-time algorithm: “deferred acceptance” (Gale and Shapley, 1962)

Stable Matching Problem with Couples (SMP-C)

- Same objective as before, but couples can apply together

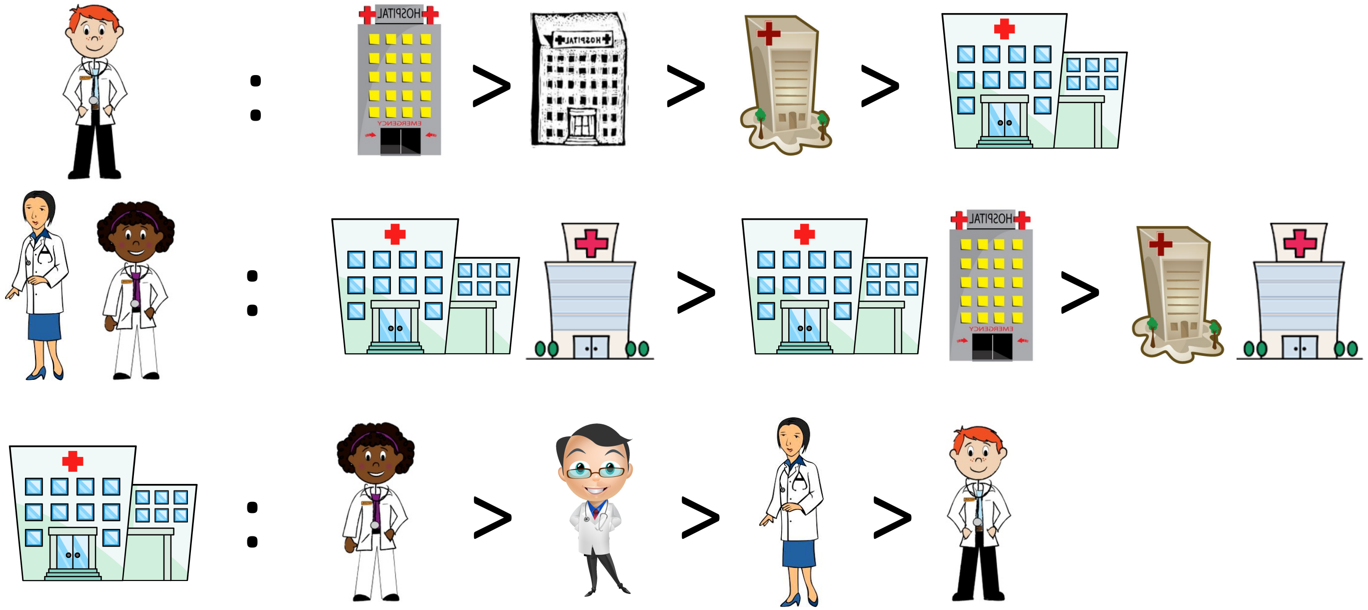

- NP-Complete

Significance of SMP-C

- United States National Resident Matching Program (*NRMP*): 34,905 residents, 6% in couples
- Smaller markets in Canada, Israel, Scotland…

THE MATCH℠
NATIONAL RESIDENT MATCHING PROGRAM®

Contributions

- Use satisfiability (SAT) encoding for SMP-C to analyze strategic properties of SMP-C
    - Analyze a conjecture and result from SMP
- Some new theory relevant to strategy-proofness in SMP-C
- Implement a mechanism for SMP-C with good strategic properties

Strategic Concerns in the NRMP

- NRMP algorithm redesigned in 90s
- New algorithm designed to make manipulation by residents as hard as possible
- Study of manipulations has focused on *truncations*

Truncation Example

Rankings

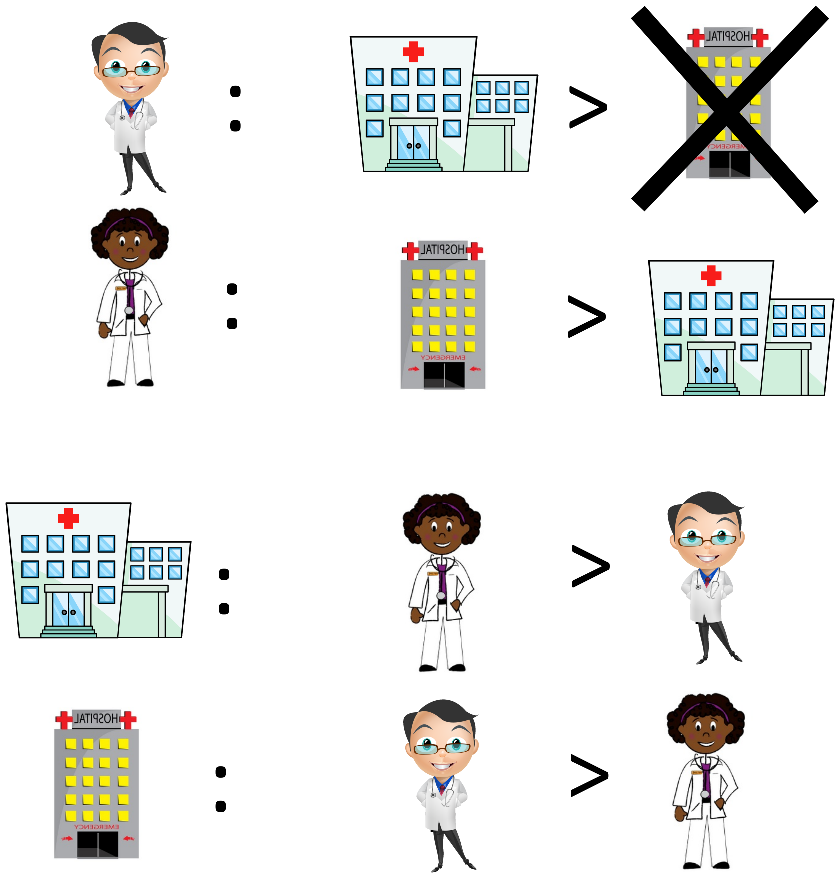

Stable Matchings

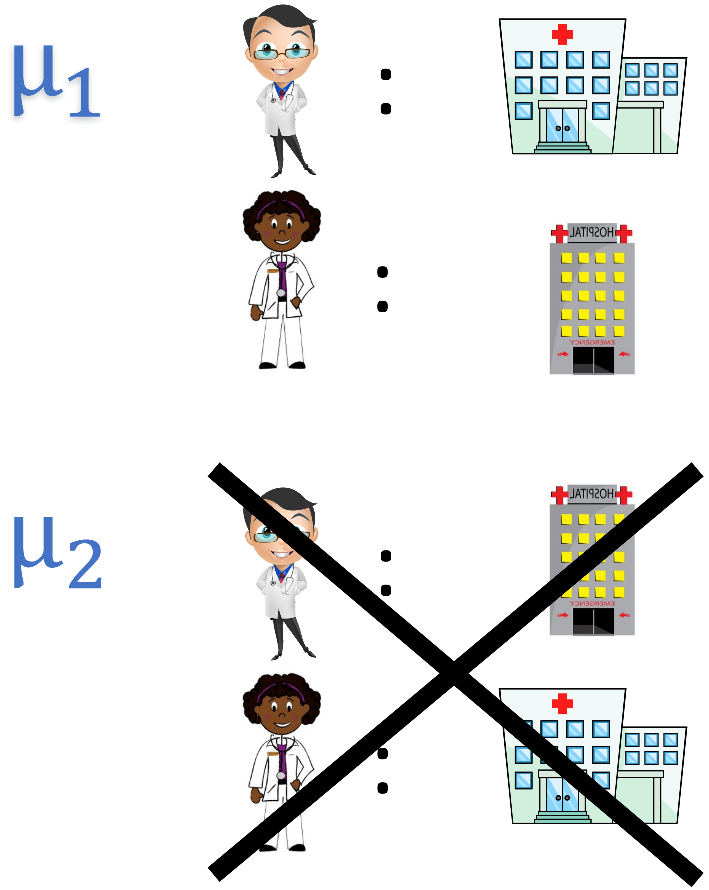

Truncations in NRMP

- Roth and Peranson (1999): at most 0.01% of residents and 0.1% of hospitals have an incentive to truncate
  - Very few opportunities for truncating on either side
  - Roth and Peranson conjectured that market size plays a role

Market Size and Strategy-Proofness in SMP

- Let n be the market size
- Let k be the preference list length
- Roth and Peranson (1999): "even when preferences are uncorrelated, as k/n becomes small, the set of stable matchings becomes small."
- Immorlica and Mahdian (2005) proved that, for SMP, expected fraction of residents with more than one stable hospital approaches zero as n approaches infinity (for fixed k)

Market Size and Strategy-Proofness in SMP

- Let *n* be the market size
- Let *k* be the preference list length
- Roth and Peranson (1999): “even when preferences are uncorrelated, as k/n becomes small, the set of stable matchings becomes small.”
- Immorlica and Mahdian (2005) proved that, for SMP, expected fraction of residents with more than one stable hospital approaches zero as *n* approaches infinity (for fixed *k*)

Outline

- Introduction and Contributions ✓
- **Theory of Strategy-Proofness**
- Finding Stable Matchings in SMP-C
- Empirical Results
- Conclusion

Why Truncations?

- In SMP, truncations are sufficient for manipulation (Roth and Vande Vate, 1991)
- Out of all manipulations, truncations can be identified with the least information about others' prefs (Roth and Rothblum, 1999)
- Easy to check empirically if a resident can benefit by truncating

Definitions: Resident Preferred

- μ is *resident preferred* ($\succcurlyeq_R$) to μ' if, for each resident or couple a, $\mu(a) \succcurlyeq_a \mu'(a)$
  - All residents and couples at least as well off

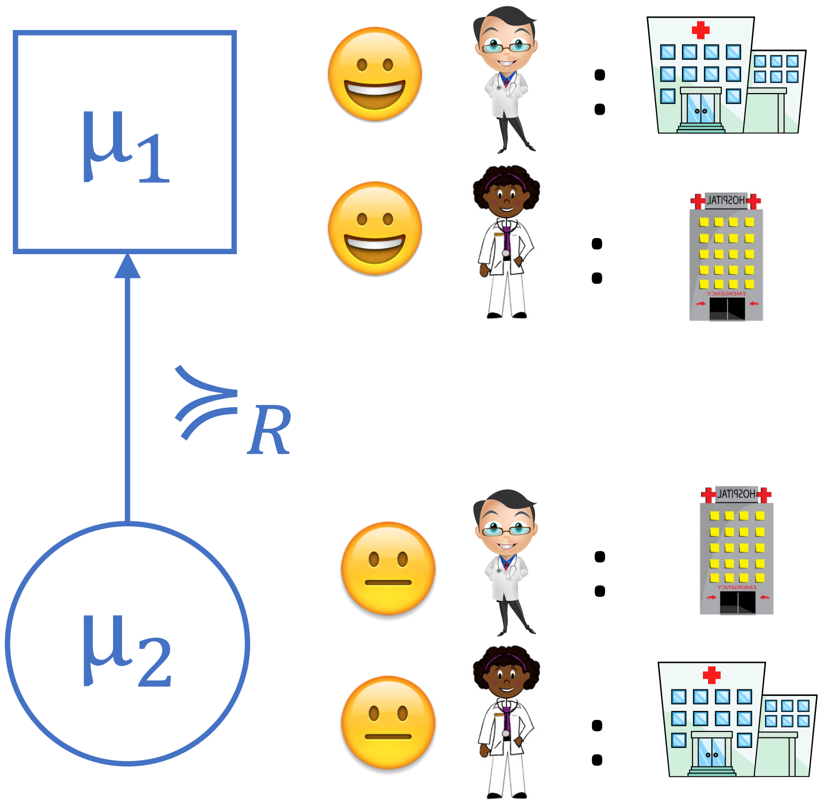

Definitions: Resident Optimal

- μ is *resident optimal* ($\mathcal{R}_{opt}$) if, for all μ', $\mu \succcurlyeq_R \mu'$
  - No resident or couple can do better in a stable matching
- Theorem (this paper): in SMP-C, residents can't benefit by truncating in an $\mathcal{R}_{opt}$ matching

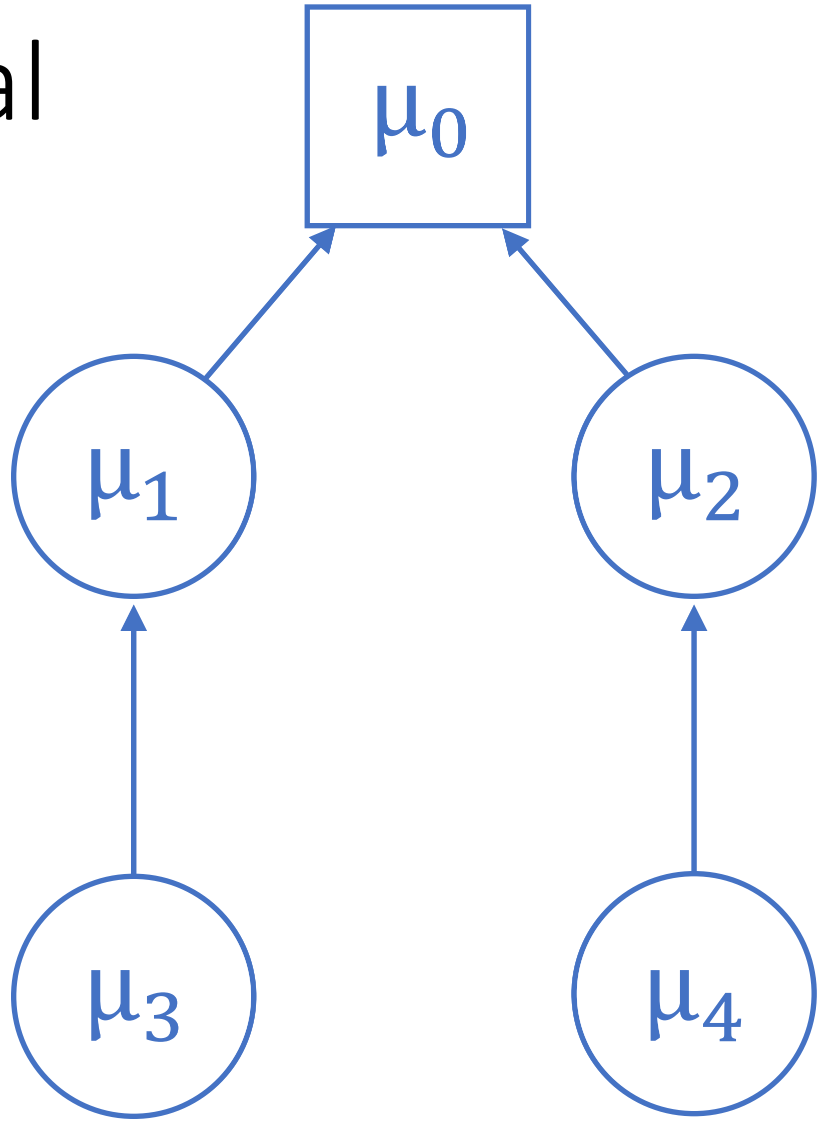

Definitions: Resident Pareto Optimal

- New, but natural extension
- μ is *resident Pareto optimal* ($\mathcal{RP}_{opt}$) if there is no μ' such that $\mu' \succcurlyeq_R \mu$
  - Always exists in SMP and SMP-C
- All $\mathcal{R}_{opt}$ matchings are $\mathcal{RP}_{opt}$

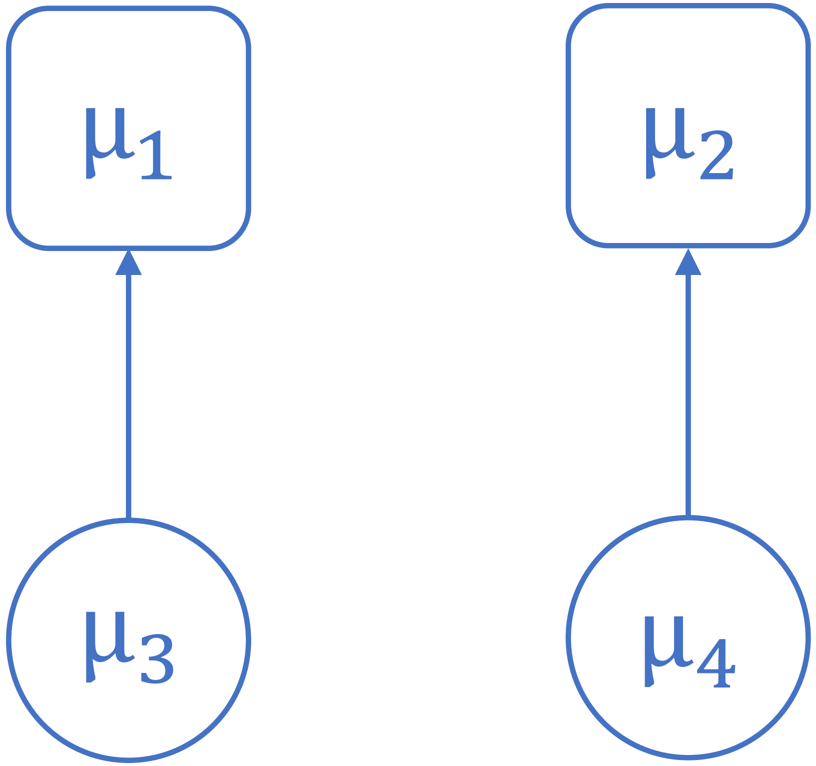

Strategy in SMP-C: Resident Pareto Optimal Matchings

- Theorem (this paper): no stable mechanism is strategy-proof against resident truncations
  - WLOG, mechanism chooses μ_1
  - Some residents prefer μ_2

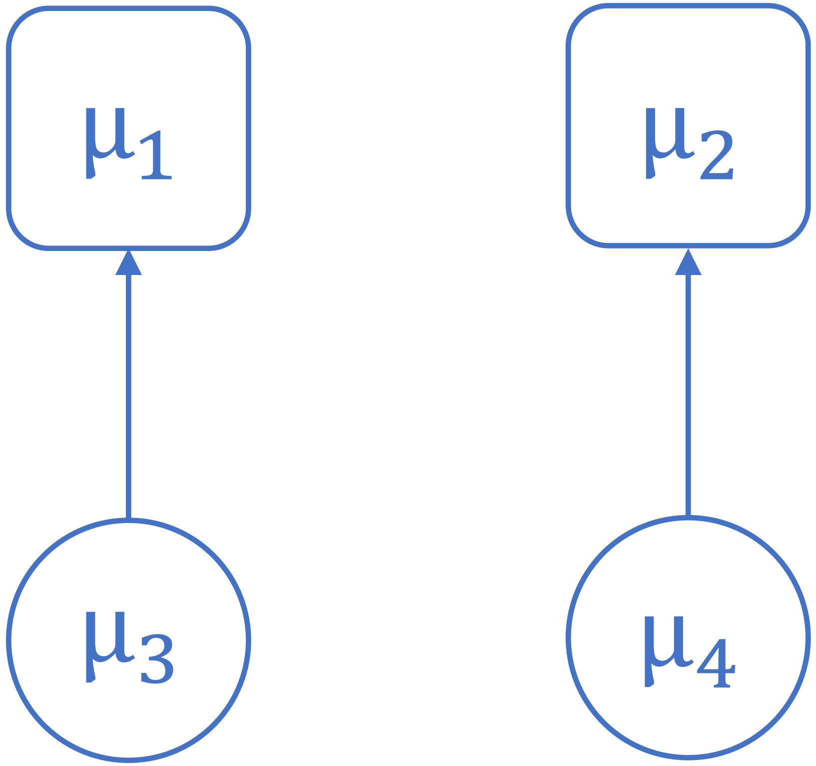

Strategy in SMP-C: Random Stable Matchings

- Theorem (this paper): random stable mechanism may be strategy-proof when $\mathcal{RP}_{opt}$ mechanism is not
  - Suppose $\mu_1 \succ_r \mu_2 \succ_r \mu_3 \succ_r \mu_4$
  - r truncates below $\mu_1(r)$
  - Truncating increases chance of being unmatched
  - Depends on utility values of ranked programs

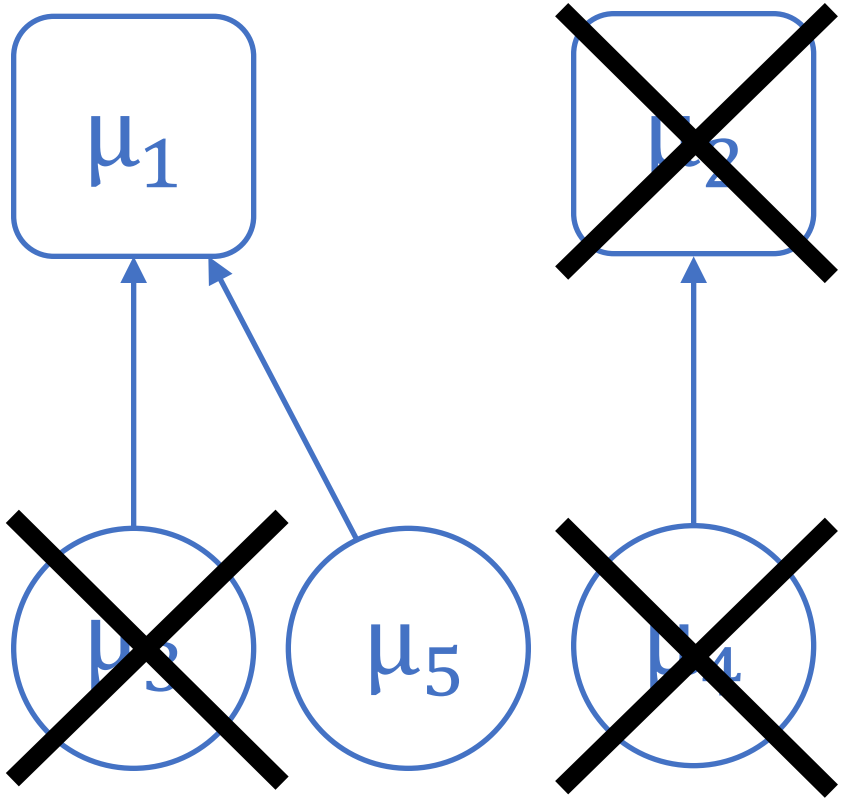

Outline

- Introduction and Contributions ✓
- Theory of Strategy-Proofness in SMP-C ✓
- **Finding Stable Matchings in SMP-C**
- Empirical Results
- Conclusion

Solvers for SMP-C

- NRMP uses “deferred acceptance” alg. (based on Gale-Shapley)
  - Relies on low % of couples (Drummond et al., 2015)
  - With low % couples, can solve large instances very fast
- Drummond et al. (2015) develop a satisfiability (SAT) encoding for SMP-C
  - Best scaling results of any complete solver

Advantages of SAT

- Can quickly find $\mathcal{RP}_{opt}$ or $\mathcal{R}_{opt}$ matchings
- Can also enumerate all stable matchings
  - Could be used to implement randomized mechanisms
- Can implement an $\mathcal{RP}_{opt}$ mechanism
  - Guaranteed to return an $\mathcal{R}_{opt}$ matching if one exists

Outline

- Introduction and Contributions ✓
- Theory of Strategy-Proofness ✓
- Finding Stable Matchings in SMP-C ✓
- **Empirical Results**
- Conclusion

Preference Models

- Impartial culture with geography (IC-GEOG) - Kojima et al. (2013)
  - Uniformly distributed (uncorrelated) preferences, couples only apply to hospitals in same region
- Scottish Foundation Allocation Scheme with geography (SFAS-GEOG) - Biró et al. (2013)
  - Geography plus Plackett-Luce
    - Hospitals and residents have varying popularity

Performance of Deferred Acceptance Algorithms

- Return $\mathcal{RP}_{opt}$ matching 90-100% of the time, i.e., 0-10% failure rate
- Also sometimes fails to find existing stable matching

Effect of Market Size

- Not affected by market size:
  - # stable matchings
  - # $\mathcal{RP}_{opt}$ matchings
  - % of instances with $\mathcal{R}_{opt}$ matching
  - % instances with at least one stable matching

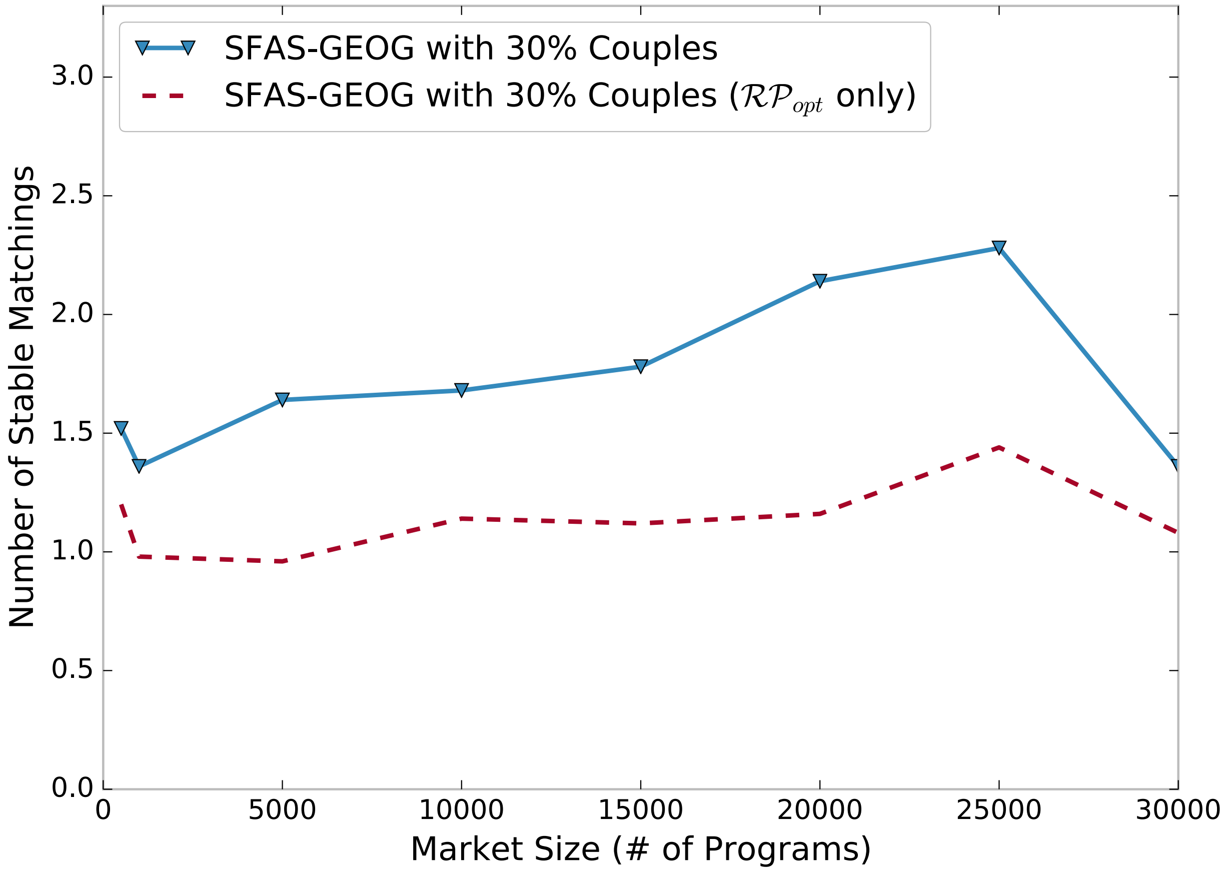

Effect of Market Size

- Not affected by market size:
  - # stable matchings
  - # $\mathcal{RP}_{opt}$ matchings
  - % of instances with $\mathcal{R}_{opt}$ matching
  - % instances with at least one stable matching

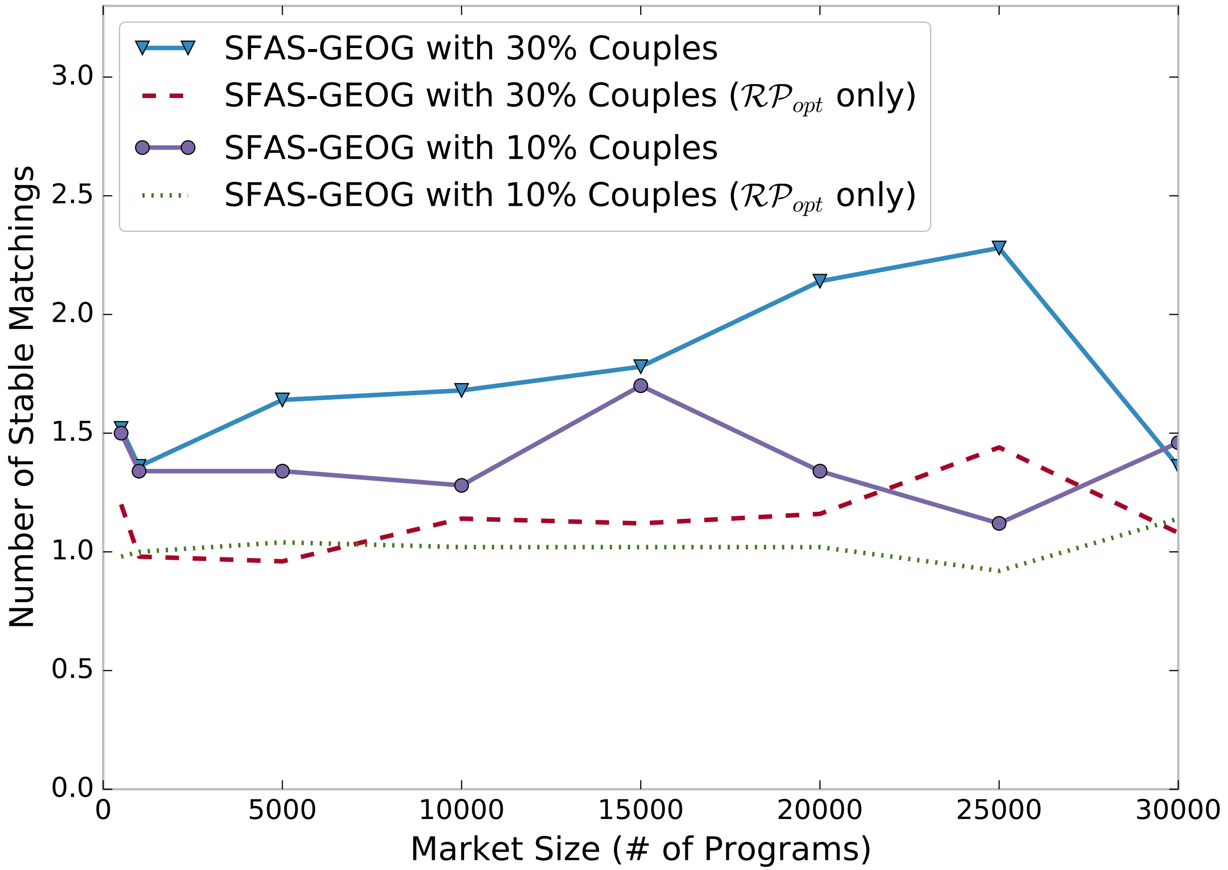

Avg. # of Add’l Stable Hospitals per Resident

- Immorlica and Mahdian’s result appears to hold for SMP-C

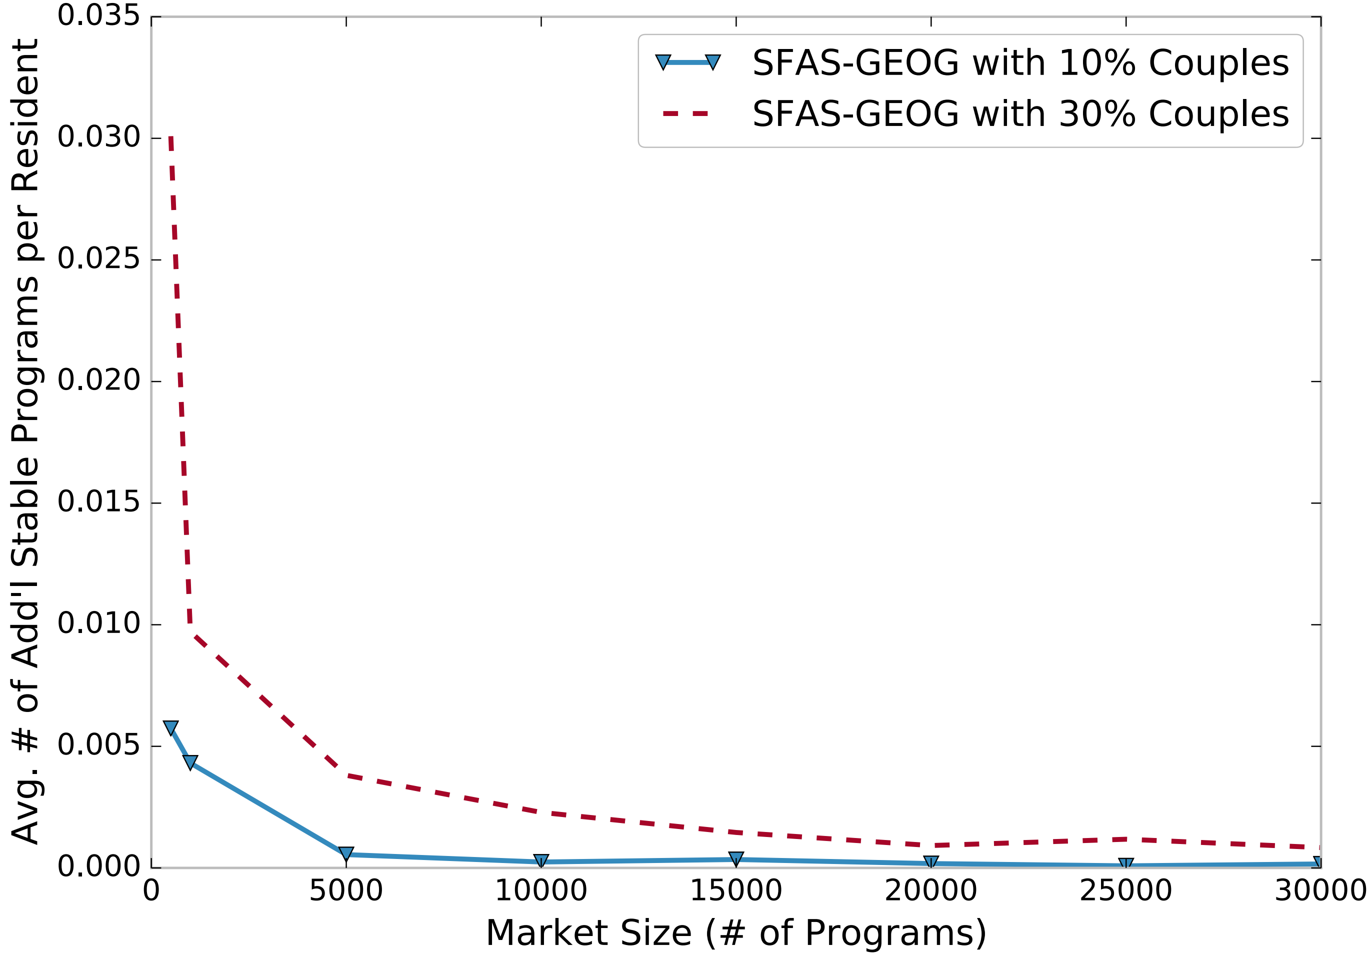

Avg. # of Residents with Incentive to Manipulate Under Truncations

- There will always be some residents with incentive to truncate under an $\mathcal{RP}_{opt}$ mechanism

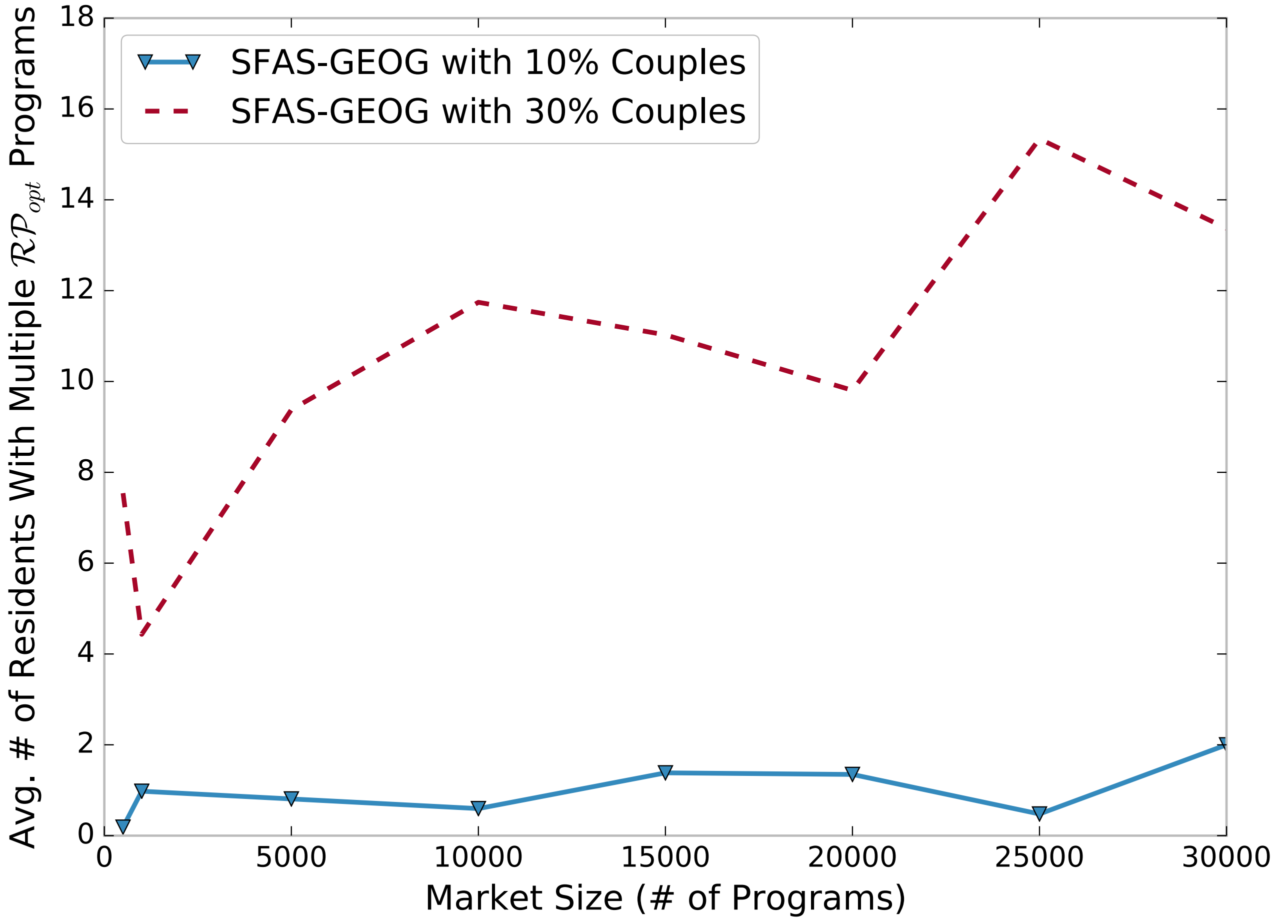

Conclusions and Future Work

- Use SAT encoding for SMP-C to show:
  - Roth and Peranson's conjecture appears false for SMP-C
  - Immorlica and Mahdian's result appears true for SMP-C
- New theory for study of strategic behavior in SMP-C
- Provide implementation of $\mathcal{RP}_{opt}$ mechanism
- Future work
  - Proofs possible?
  - Study more general class of manipulations—reorderings
  - Use of randomization for greater strategy-proofness

Thank You! Questions?

- Poster tomorrow (Thursday)
- Code available online at git.io/vwlXq or link from website

A Caveat: Reorderings

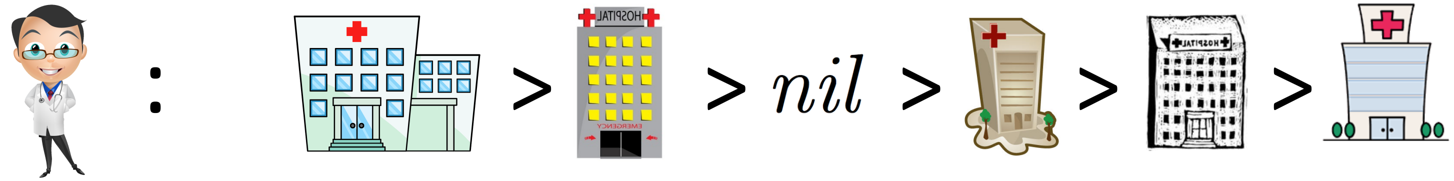

- Informally, for truncations, only need to look at set of stable matchings under true preferences (Roth and Vande Vate, 1991) (analogue for SMP-C proved in this paper)
- In SMP-C, *reorderings* can create stable matchings that are not stable under true preferences (Biró and Klijn, 2011)
- Reorderings hard to analyze computationally
  - May also be hard for manipulators to find

Add’l RP_{OPT} Hosptials, % of Instances with Stable Matching

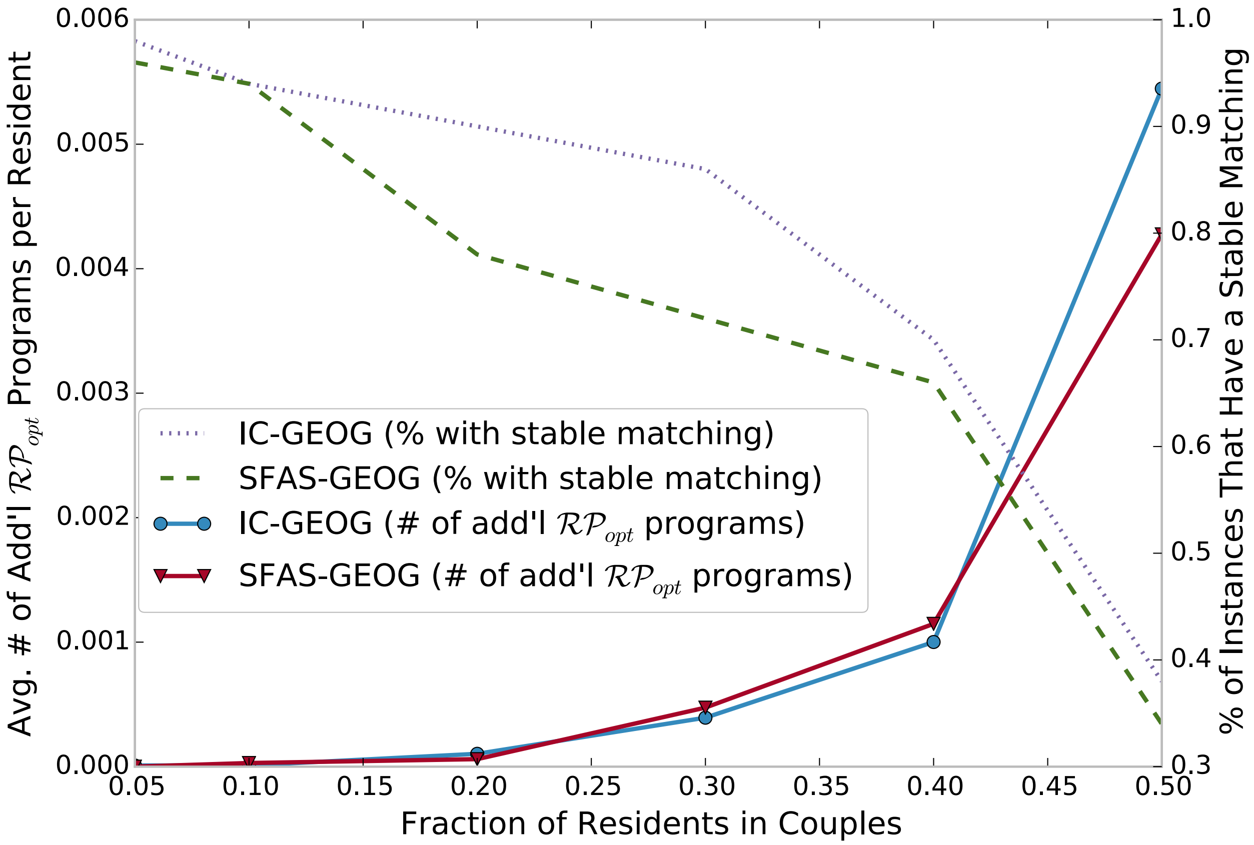

Resident Optimal Matchings as Market Size Grows

- Not affected by market size varying between 250 and 30,000 residents
  - R_{OPT} exists 90-95% with 10% couples
  - R_{OPT} exists 60-70% with 30% couples
- TODO: insert graph

Definitions: Resident Pareto Optimal

- μ is *resident optimal* (R_{OPT}) if, for all μ', $\mu \succcurlyeq_R \mu'$
  - Always exist for SMP, but not SMP-C
- Theorem (this paper): in SMP-C, residents can't benefit by truncating in an R_{OPT} matching