Fairness in Matching under Uncertainty

Siddartha Devic (USC), David Kempe (USC), Vatsal Sharan (USC), Aleksandra Korolova (Princeton)
@sid_devic, @korolova

https://arxiv.org/abs/2302.03810

Motivation: Two-Sided Marketplaces

- ***Two-sided marketplaces*** are increasingly relevant
  - e.g., assigning students to schools, users to advertisers, applicants to job interviews, etc.
- Participants in the marketplace:
  - Individuals (ride-hailers, job candidates, social media users)
  - Resources (drivers, jobs, ads)
- How can we consider and optimize fairness desiderata in these complex systems, often in conflict with utility?

Intuition for Fairness Axiomatization

- Two individuals, Alice and Bob, both submit their CVs.
- How do we hire fairly?

Individual fairness [3] would have us allocate:

Alloc. Diff. () ≤ distance(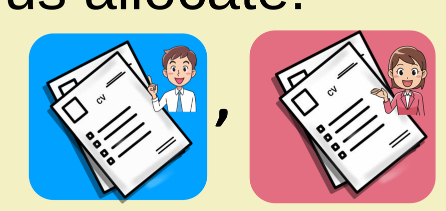)

Singh, Joachims, Kempe (2021) [1] propose a randomized approach which utilizes uncertainty as a **cornerstone of fairness**

Observables **Posterior distributions for merit**

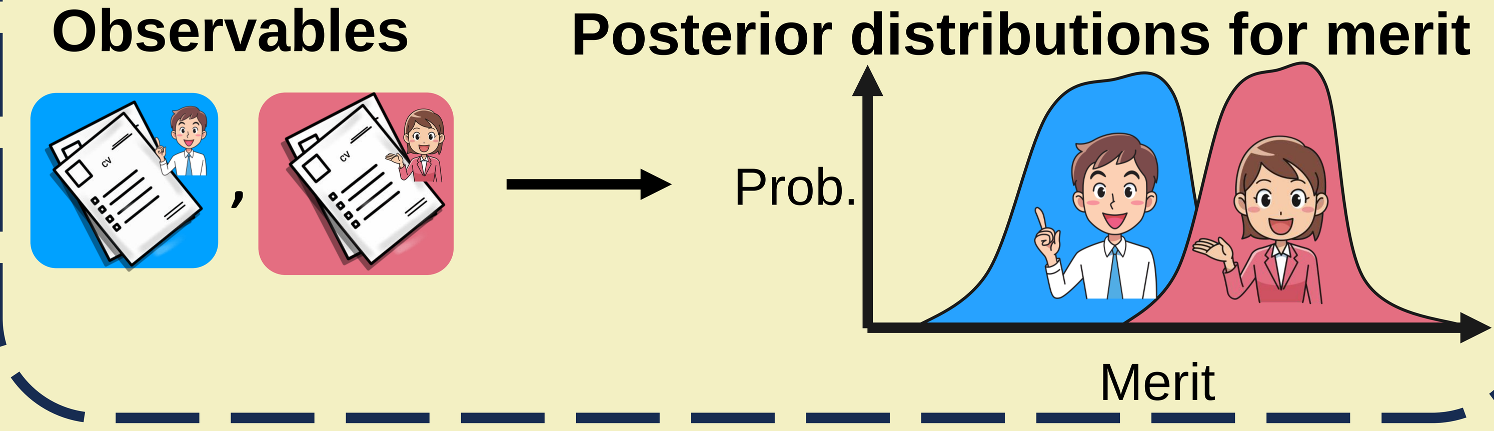

- Importantly, **ML algorithms** often output **distributions** over merit, scores, fitness, etc.

Axioms for Fair Decision Making

- **Axiom 1** (Full information, no uncertainty): Always pick the candidate with greater merit.
- **Axiom 2** (Uncertainty present): Make randomized decision proportional to possible futures implied by uncertainty.

An Approach Based on Possible Futures

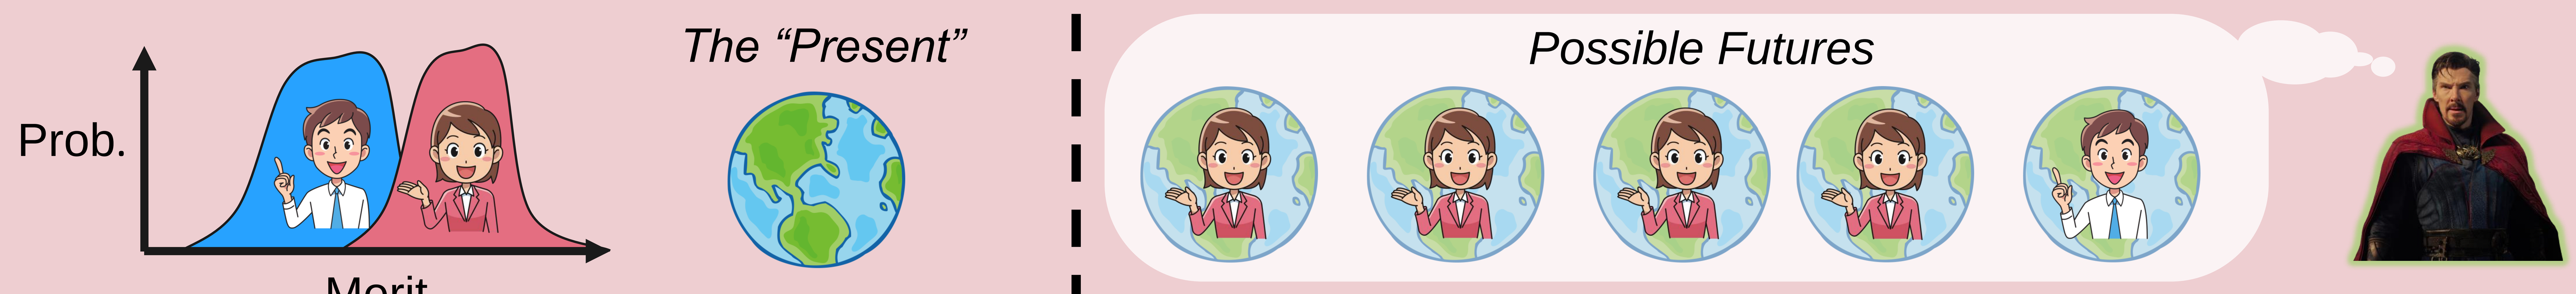

(Alice ≥ Bob w.p. 0.8) (Which of Alice / Bob is more qualified in each of these futures)

- *Singh et al.*: Prob. to select Alice / Bob in **present** ≥ prob. Alice / Bob more qualified in **possible futures** (select .80 / .20)

Tradeoffs Between Utility and Fairness

- **Axiom 2** is often in conflict with utility: in hiring, one utility maximizing solution is to always select Alice
  - In contrast, fair solution selects Alice w.p. 0.8
- We allow a multiplicative *relaxation* of fairness in order to tradeoff with utility: call this **ϕ-fair** for $\phi \in [0,1]$.

Prob. Alice hired ≥ ϕ · Prob. Alice more qualified in possible futures

Prob. Bob hired ≥ ϕ · Prob. Bob more qualified in possible futures

Application: Two-sided Marketplaces

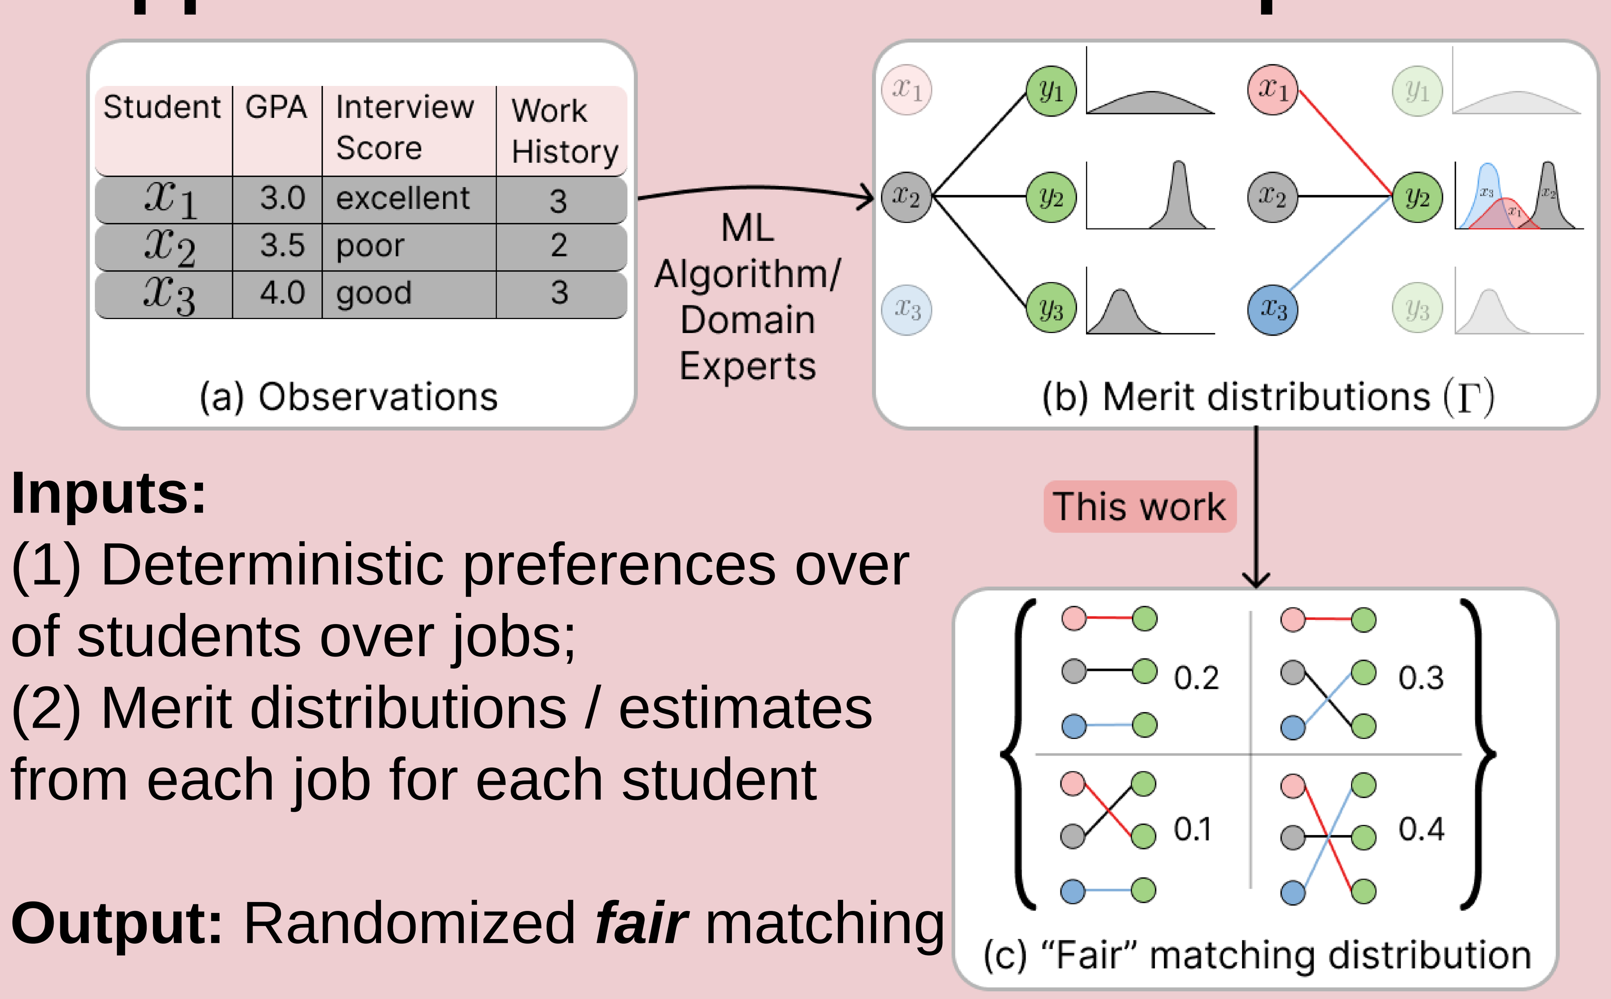

Inputs:
(1) Deterministic preferences over of students over jobs;
(2) Merit distributions / estimates from each job for each student

Output: Randomized ***fair*** matching

Fairness is w.r.t. to "possible futures" where we *sample* merits for each candidate and compute a **stable** match.

References [1]: A. Singh, D. Kempe, and T. Joachims. Fairness in ranking under uncertainty. In Proc. 35th Advances in Neural Information Processing Systems, pages 11896–11908, 2021. [2]: Lukáš Brožovský and Václav Petříček. Recommender system for online dating service. In *Proc. Znalosti*, pages 29–40, 2007. [3] C. Dwork, M. Hardt, T. Pitassi, O. Reingold, and R. S. Zemel. Fairness through awareness. In Proc. 3rd Innovations in Theoretical Computer Science, pages 214–226. ACM, 2012

Theoretical Results

- Run a **Linear Program** to maximize utility for $\phi \in [0,1]$.
- However, we can only **estimate** the distribution over matchings in possible futures through sampling.

Theorem 1 (Informal): Running our LP using a *sampled* ϵ close estimate of possible futures achieves $\left(\frac{1}{\phi n\epsilon+1}\right)$-approx. of max utility and $\left(\frac{\phi\left(1+\frac{\epsilon}{2}\right)}{n\epsilon+1}\right)$-fairness relative to the LP run with the true possible future distribution.

Theorem 2 (Informal): The analysis of our method is tight for both fairness and utility.

Empirical Result

- **Ours** vs. **Thompson sampling baseline** on [2]. We observe a utility gain for all levels of ϕ, even at $\phi = 1$ (full fairness)!

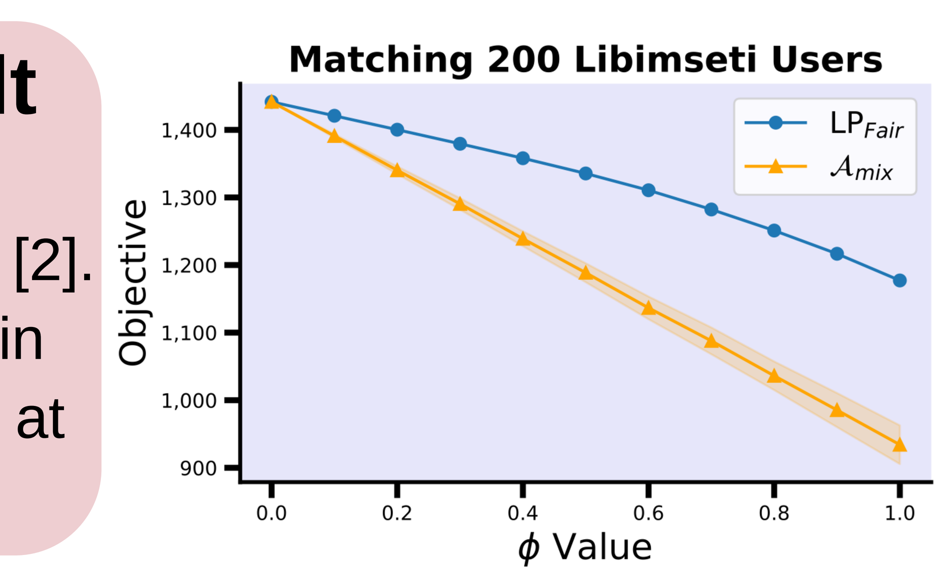

Key Results

- Axiomatize a notion of ***individual fairness*** in two-sided marketplaces which respects the uncertainty in the merits.
- Design a **linear programming** framework to find fair utility-maximizing distributions over allocations.
- Prove that LP is **robust to approximate estimations** of the uncertain merit distributions, a key property in combining the approach with ML techniques.
- Verify the method empirically by designing an experiment in a two-sided market derived from a dating app.