Fairness in Matching under Uncertainty

NDSEG Fellows Conference, 2023

Siddartha Devic (USC, devic@usc.edu)
David Kempe (USC)
Vatsal Sharan (USC)
Aleksandra Korolova (Princeton)

Two-sided matching markets

- Marketplaces are becoming increasingly relevant
  - Ride-hailing, job matching, ad serving
- 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?

Intuition: A simple hiring setting

- Two individuals, Alice and Bob
- Trying to select one person for a job, both Alice and Bob submit a CV.

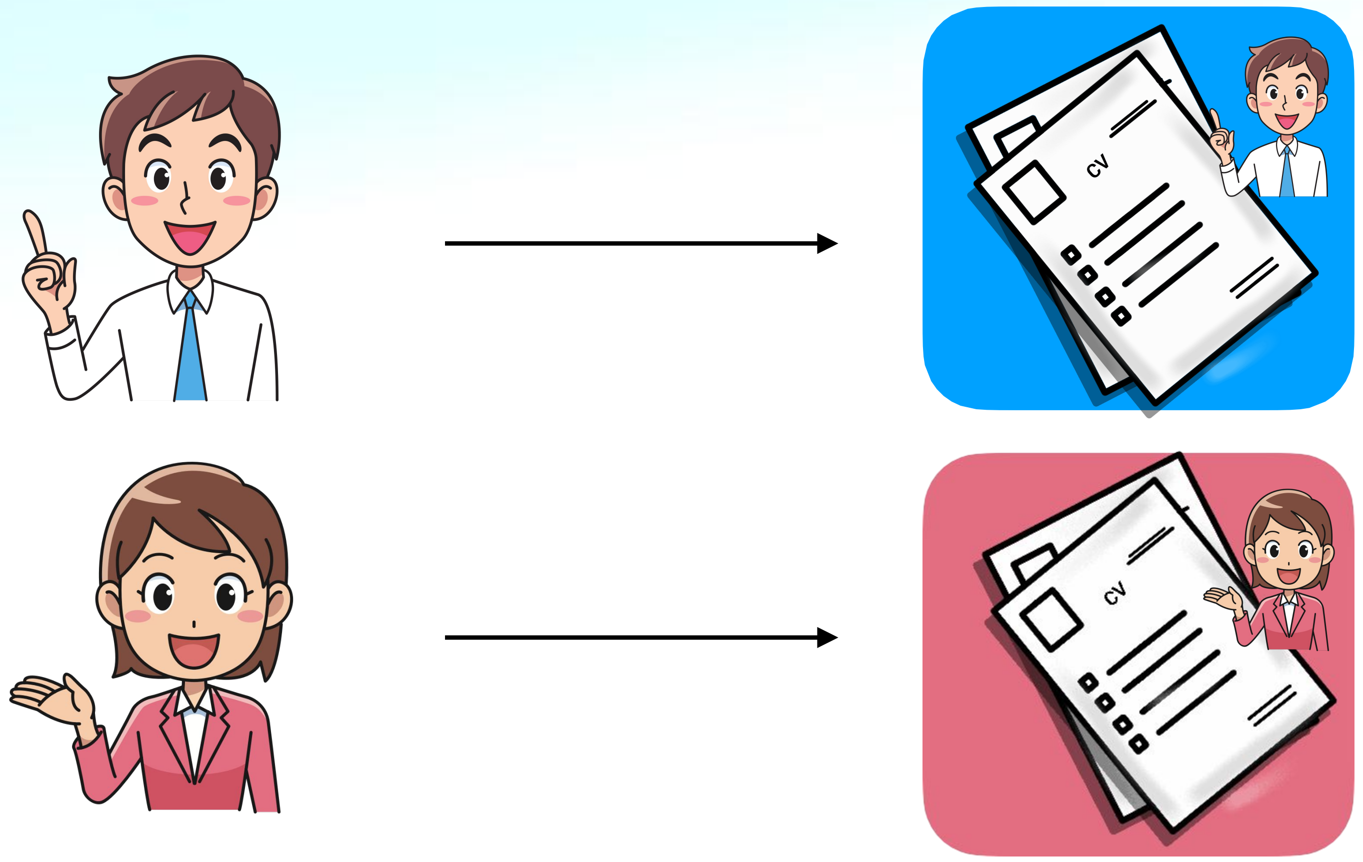

- How do we make a fair hiring decision with the CVs in hand?

Traditional individual fairness approaches

- Individual fairness (Dwork, Hardt, Pitassi, Reingold, Zemel 2012):

Allocation diff. ( **,**  **)** $\leq$ **distance (**  **,**  **)**

- For some distance metric on CVs.

Uncertainty as the cornerstone of fairness

- Singh, Joachims, Kempe (2021) propose a randomized approach which utilizes uncertainty

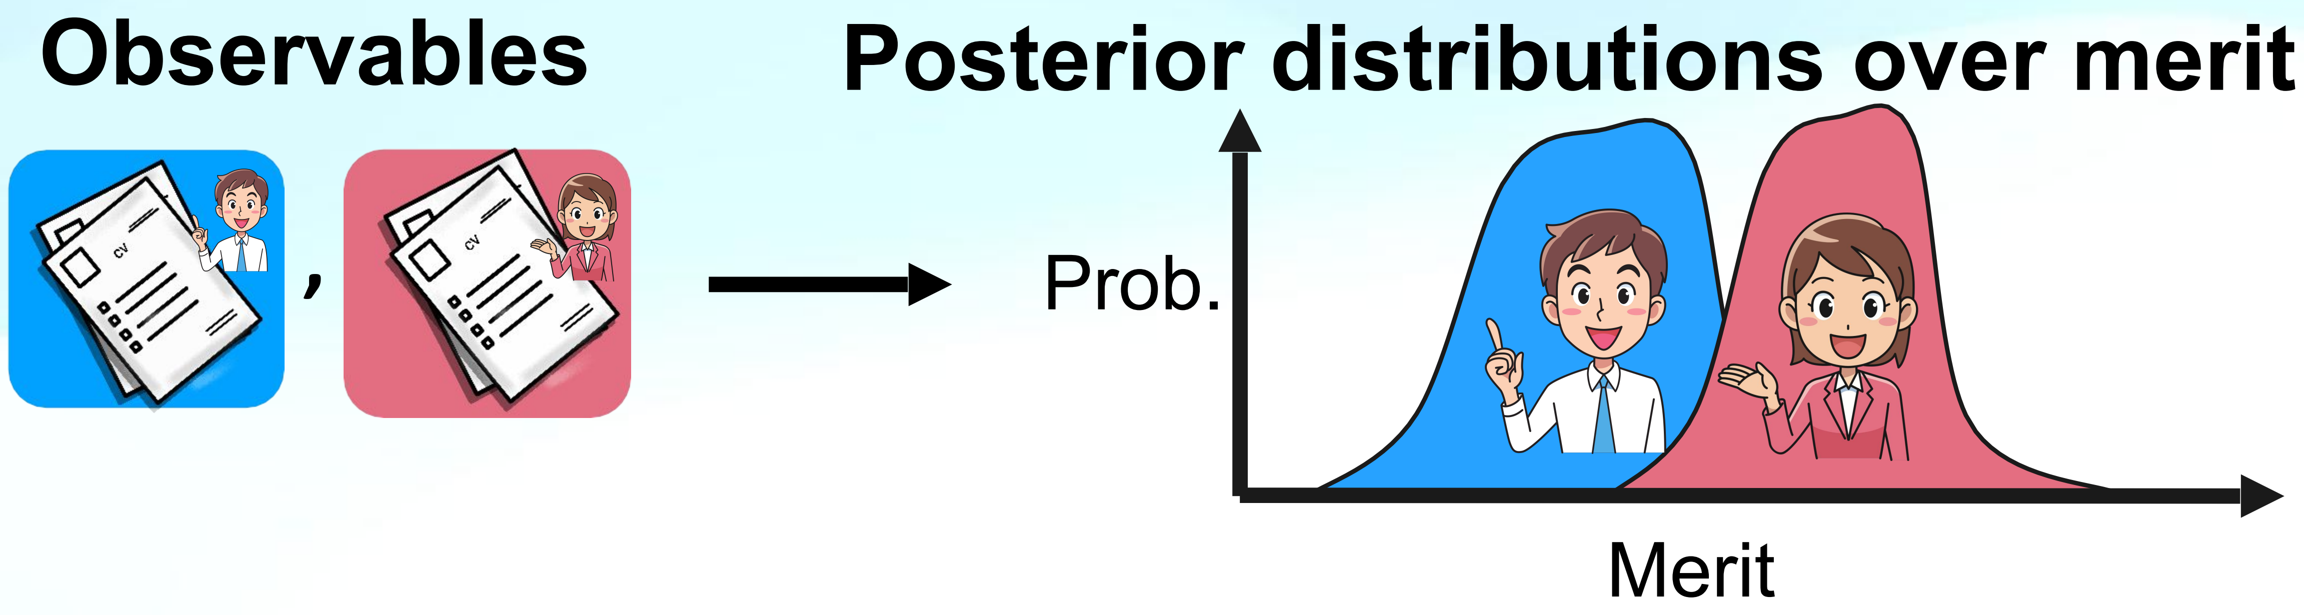

- Importantly, **ML algorithms** often output **distributions** over merit

An approach based on possible futures

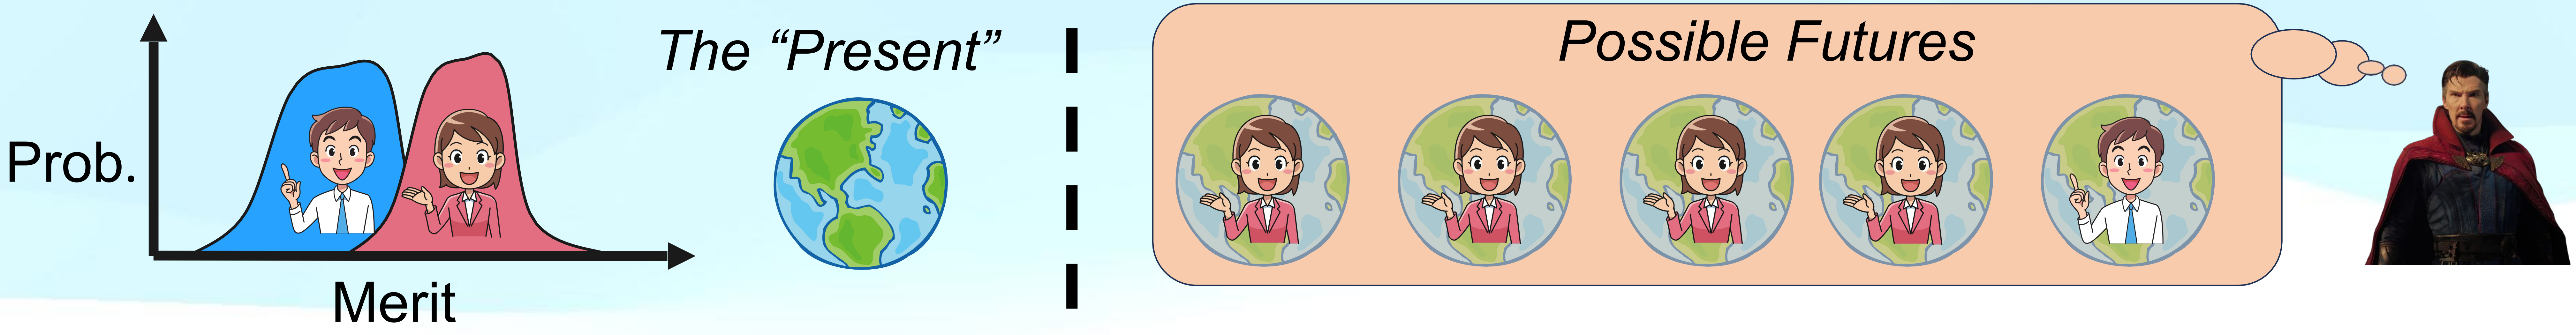

(**Alice** $\geq$ **Bob** w.p. 0.8)

(Which of **Alice** / **Bob** is more qualified in each future)

- Singh et al. [1]: Prob. to select **Alice** / **Bob** in **present** ≥ Prob. **Alice** / **Bob** is more qualified in **possible futures**
  - e.g., select with **80** / **20**

Generalizing the axiomatic framework of Singh et al.

- **Axiom 1** (No uncertainty, meritocracy): We should always pick the candidate with greater merit.
  - If we were in the future where **Alice** was more qualified, we should always select **Alice**.

- **Axiom 2** (Uncertainty present): We should respect the possible futures it implies, and make a (randomized) decision proportional to these futures.
  - We don't know which future we will go to, so select proportionally (80/20).

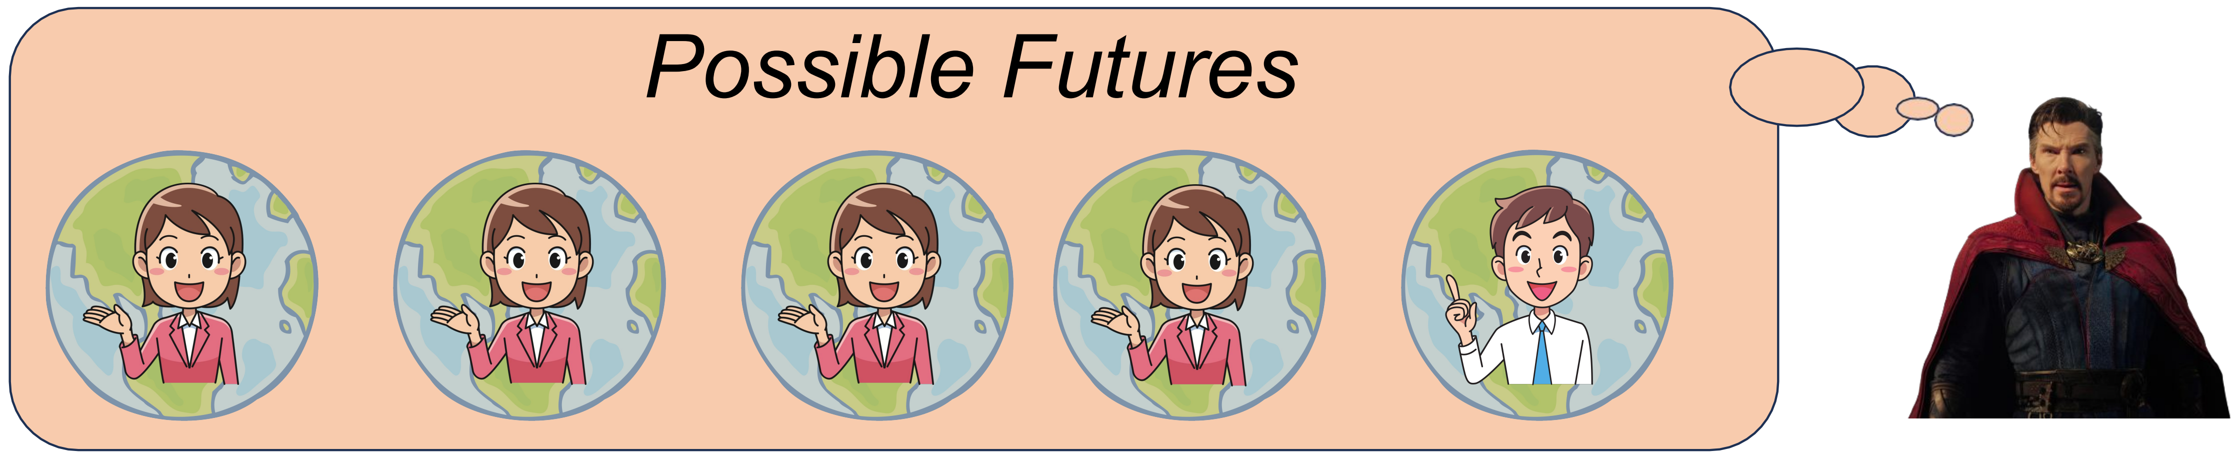

Tradeoffs Between Utility and Fairness

- **Axiom 2** is often in conflict with utility: in hiring, one utility maximizing solution is to always select **Alice**
  - This is in contrast with the fair solution, which selects **Alice** w.p. **0.8**

- We allow a multiplicative relaxation of fairness to tradeoff utility by a parameter $\phi \in [0,1]$.

- We call this ϕ-fairness.

Prob. **Alice** hired $\geq \phi \cdot$ Prob. **Alice** more qualified in possible futures

Prob. **Bob** hired $\geq \phi \cdot$ Prob. **Bob** more qualified in possible futures

An Instantiation: Two-sided Marketplaces

Inputs:

(1) Deterministic preferences of students;

(2) Merit distributions / estimates from each job for each student

Output: Randomized fair matching

Fairness is with respect to “possible futures” where we *sample* merits for each candidate compute a **stable** match.

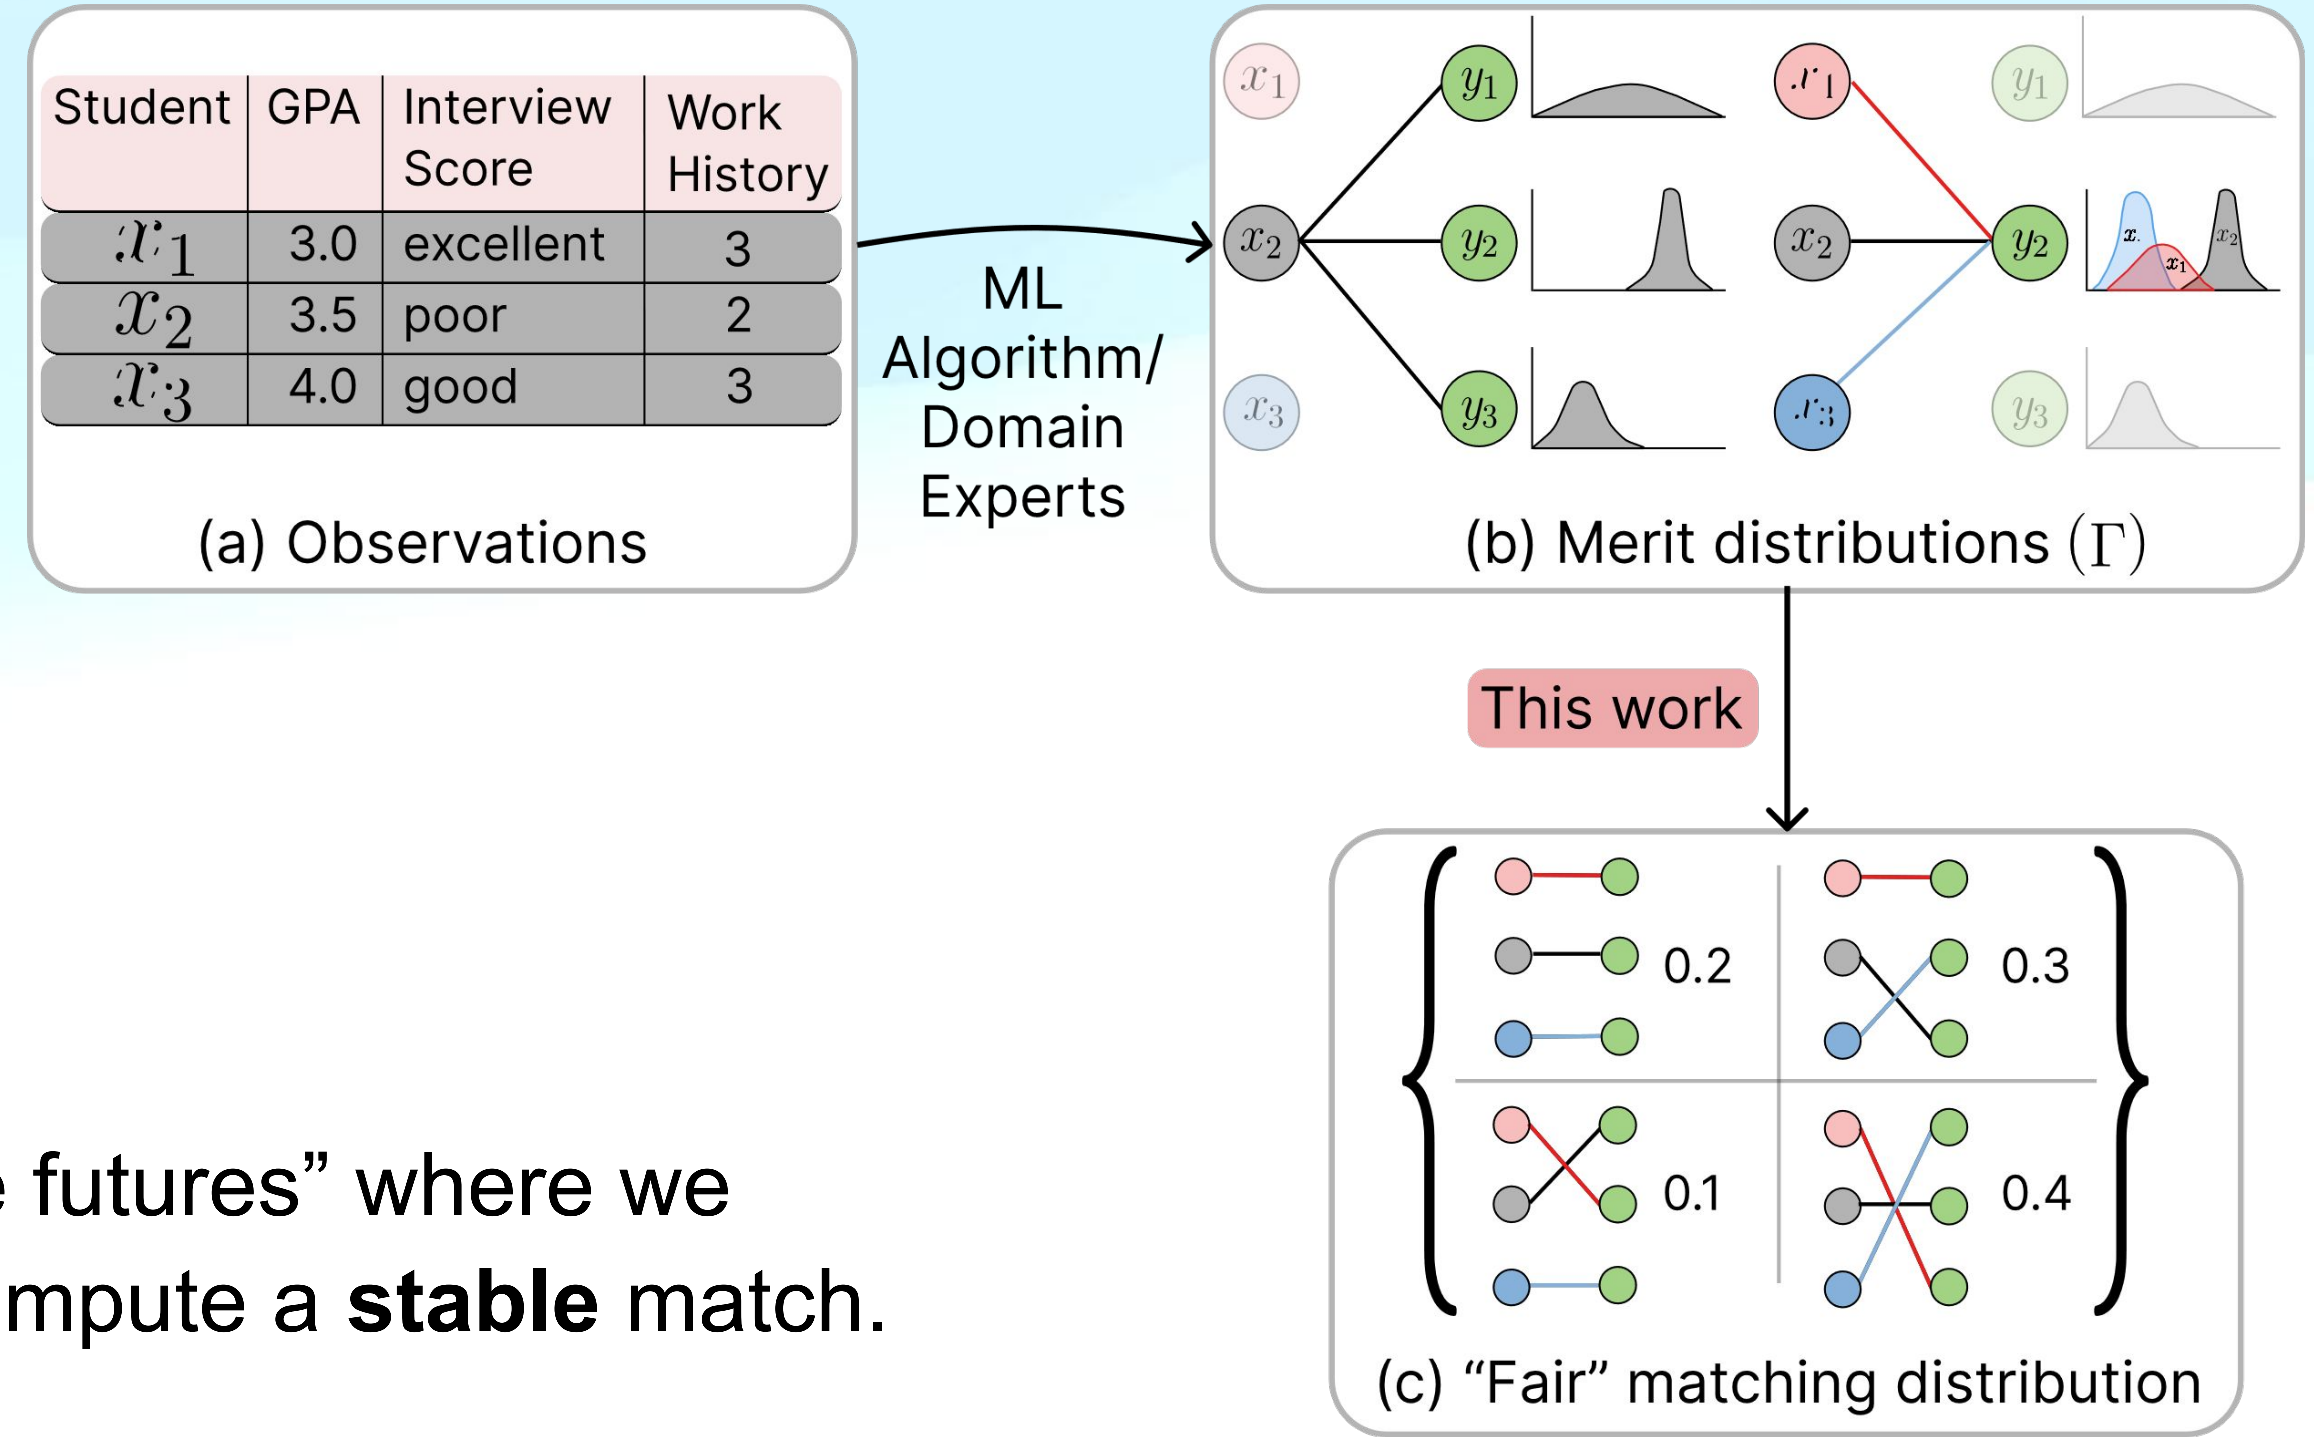

Theoretical Results

- Run a **Linear Program** to achieve max utility for relaxed fairness constraint
- However, we can only **estimate** the distribution over matchings in possible futures through sampling.

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

Theoretical Results

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

Theorem 2 (Informal): This analysis is the best we can hope for using our method.

Empirical Result

- Libimseti dating site dataset [2]. 100 users on each side of the market, used a matrix completion technique to get estimated distribution over merits.

- **Ours** vs. **Thompson sampling baseline**. Util gain even at ϕ = 1 (full fairness)!

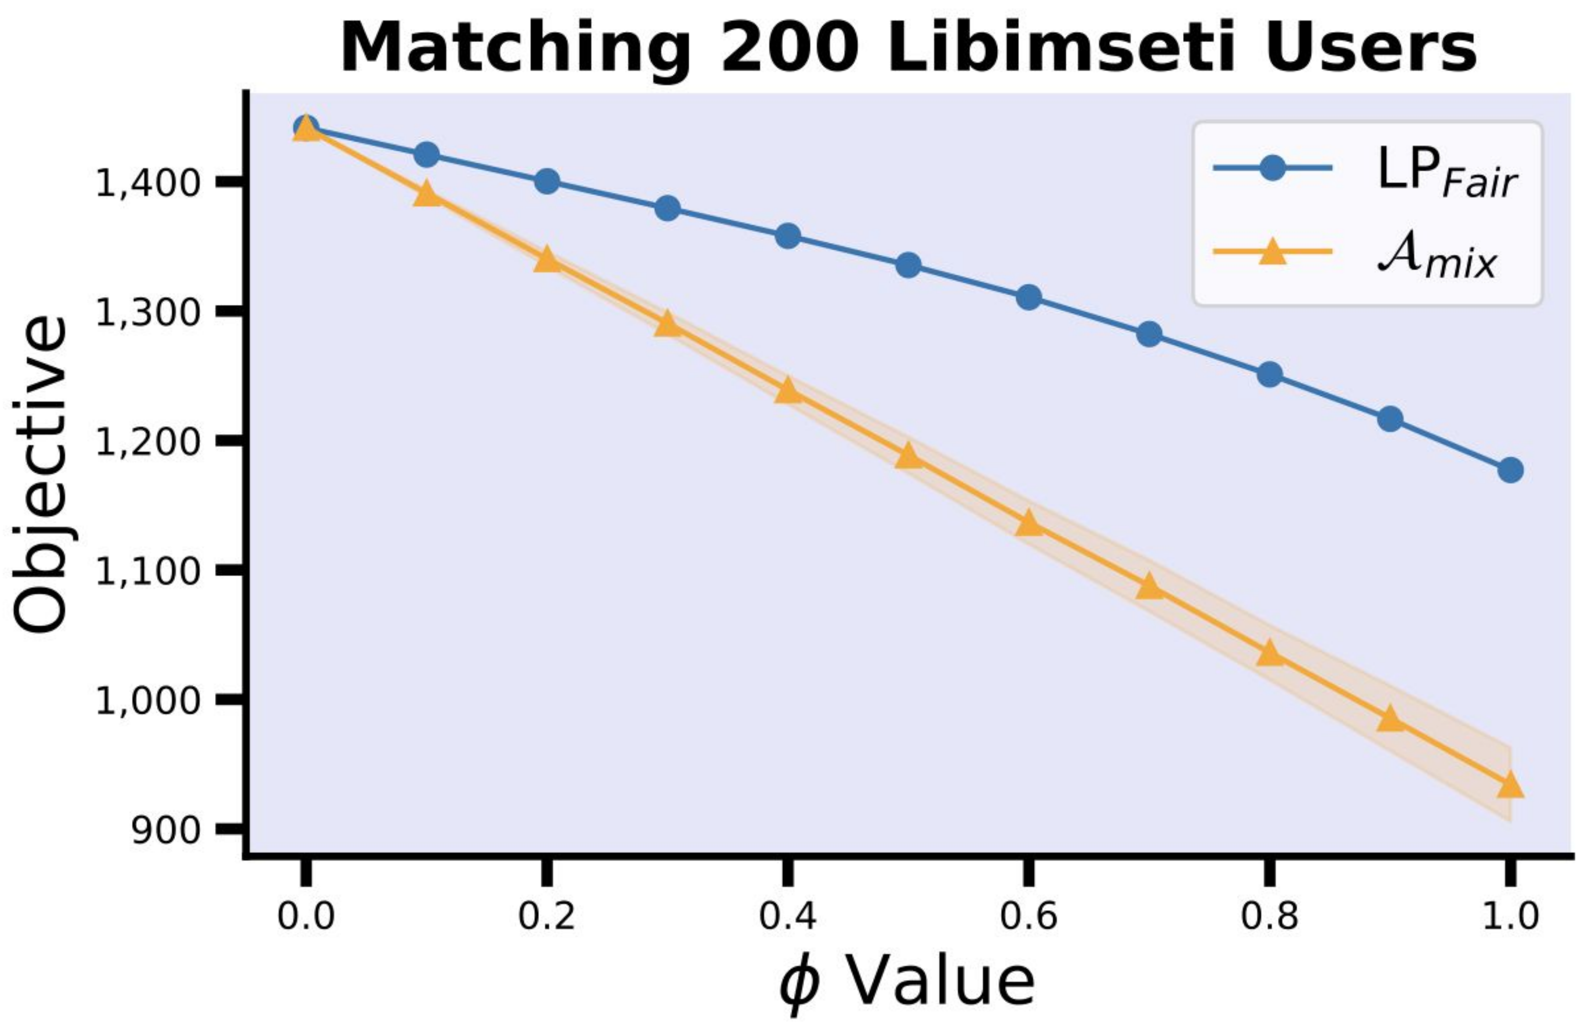

Potential Applications to DOD

- Matching problems are ubiquitous
- Assigning graduating midshipmen from the US Naval Academy to positions in the Navy
    - Naval Specialities (Surface warfare officer, Intel, etc.) have preferences over graduating midshipmen
    - Midshipmen have preferences over special
- Our framework ensures that the navy derives maximal utility from this matching subject to respecting the preferences of the midshipmen (when there is uncertainty about how they may perform on the job)

Key Takeaways

- 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 the LP is robust to perturbations in the estimated parameters of the uncertain merit distributions, a key property in **combining the approach with ML techniques.**
- Verify our method empirically by designing and running an experiment in a two-sided market with data from a dating app.
- Thank you! Questions?