Do Labor Market Frictions Reduce Employment in Low-Income Countries? Evidence from Ethiopia

Stefano Caria Pieter Gautier Lukas Hensel
Francesco Losma Marc Witte

May 12, 2026

HiJoS Workshop

Outline

1. **What constraints labor market matching in LICs?**

2. **How can we raise total employment and match quality?**
    2.1 Certification and tests
    2.2 Elicitation of preferences
    2.3 Recommendation algorithm

3. **How can we identify total employment and match-quality effects?**

- We have robust evidence that reducing job-search frictions improves the job finding rate of treated jobseekers (with respect to controls) in LMICs
- And new evidence that reducing recruitment frictions leads treated firms to fill more vacancies than the controls.

- We have robust evidence that reducing job-search frictions improves the job finding rate of treated jobseekers (with respect to controls) in LMICs
- And new evidence that reducing recruitment frictions leads treated firms to fill more vacancies than the controls.

- But what is the impact of these frictions on *total* employment and match quality?
- This is hard to know:
  - One-sided experimental designs do not quantify displacement
  - Estimands often do not map to theoretical constructs, making structural estimation hard.

- Key question in LICs, where increasing wage employment is an urgent policy priority.

Outline

1. **What constraints labor market matching in LICs?**

2. **How can we raise total employment and match quality?**
   - 2.1 Certification and tests
   - 2.2 Elicitation of preferences
   - 2.3 Recommendation algorithm

3. **How can we identify total employment and match-quality effects?**

Fundamental challenge: Vast majority of applicants are considered a poor match by firms.

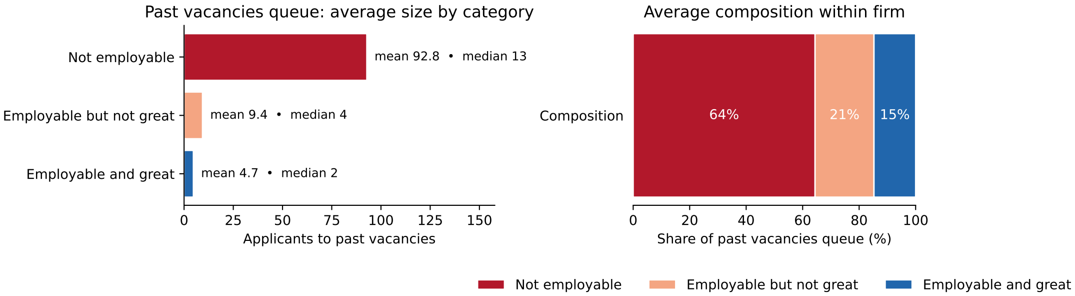

Some vacancies are not filled, some are filled with poorly matched candidates

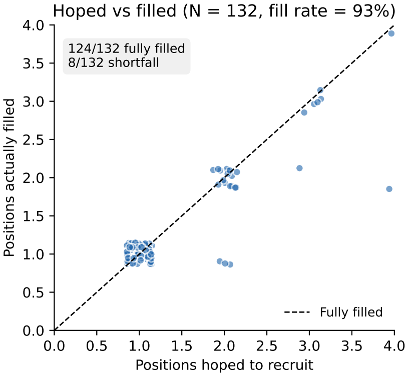

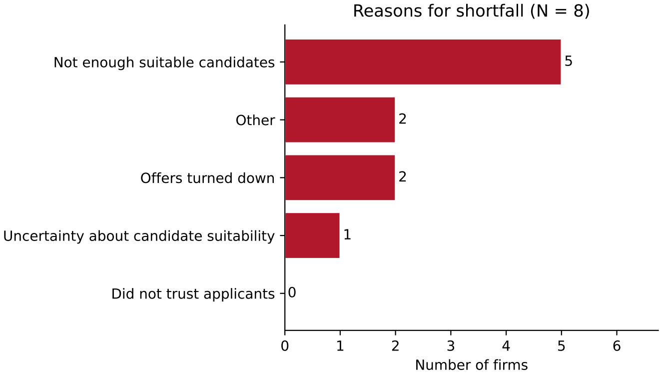

There seem to be several good matches available

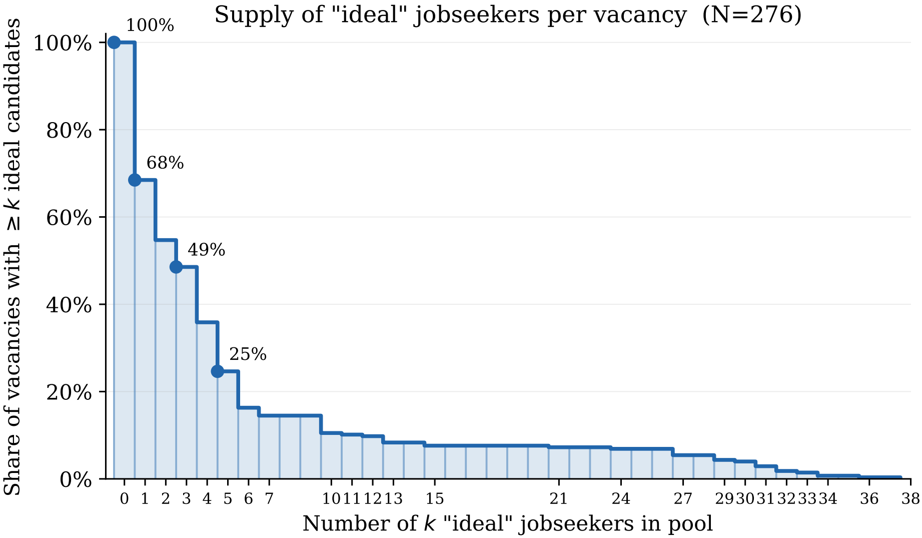

Do poor match outcomes arise because good applicants are unevenly distributed? (search location friction)

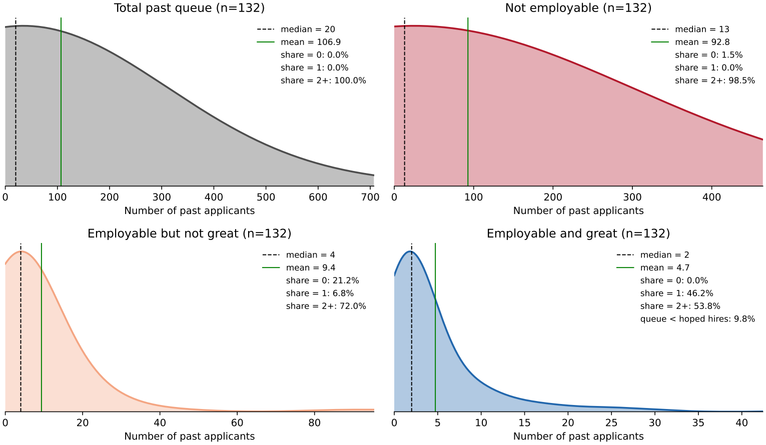

Additional applicant requests

Or because firms are unable to identify the ability of all their applicants? (screening friction)

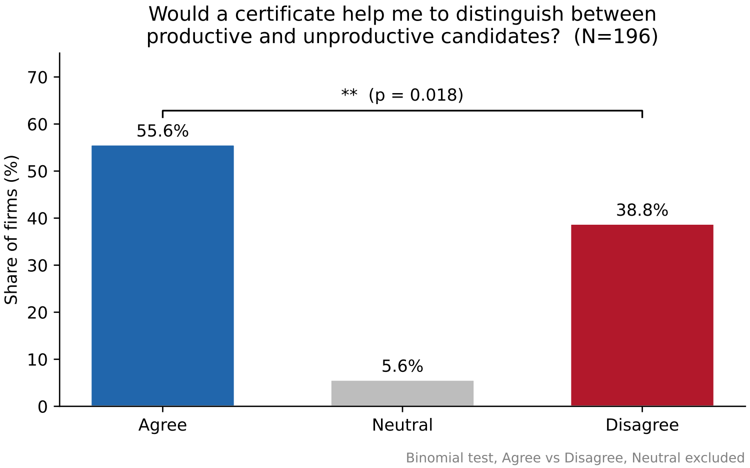

A model of location and screening frictions

- Recent labor market models can be used to parsimoniously integrate both frictions in a single framework (Cai, Gautier, and Wolthoff, 2023; Cai, Gautier, and Wolthoff, 2025)

A model of location and screening frictions

- Recent labor market models can be used to parsimoniously integrate both frictions in a single framework (Cai, Gautier, and Wolthoff, 2023; Cai, Gautier, and Wolthoff, 2025)
- Assume that workers are characterised by an ability type $x \in \mathcal{X}$. For simplicity, assume that $\mathcal{X} = \{x_L, x_H\}$ for low (employable but not great) and high ability (employable and great) types, with z the fraction of low ability type workers

A model of location and screening frictions

- Recent labor market models can be used to parsimoniously integrate both frictions in a single framework (Cai, Gautier, and Wolthoff, 2023; Cai, Gautier, and Wolthoff, 2025)
- Assume that workers are characterised by an ability type $x \in \mathcal{X}$. For simplicity, assume that $\mathcal{X} = \{x_L, x_H\}$ for low (employable but not great) and high ability (employable and great) types, with z the fraction of low ability type workers
- Firms are characterised by a type $y \in \mathcal{Y} = [\underline{y}, \bar{y}]$, and the measure of firms with types weakly below y is $J(y)$ with $J(\bar{y})$ normalised to 1

A random search process

- Under random search, a submarket $\Omega(y)$ consists of all the firms of type y and all the workers (of both types) who are exogenously assigned to this market
- Each worker places one application within the submarket.

A random search process

- Under random search, a submarket $\Omega(y)$ consists of all the firms of type y and all the workers (of both types) who are exogenously assigned to this market
- Each worker places one application within the submarket.

Define:

- Expected queue length $\lambda(y)$: the ratio between the total number of employable workers and the number of firms of type y.

A random search process

- Under random search, a submarket $\Omega(y)$ consists of all the firms of type y and all the workers (of both types) who are exogenously assigned to this market
- Each worker places one application within the submarket.

Define:

- Expected queue length $\lambda(y)$: the ratio between the total number of employable workers and the number of firms of type y.
- Expected queue length of great applicants $\mu(y)$: the ratio between the total number of workers of high type x_H and the number of firms of type y.

Modelling the search location friction

- Microfoundation: Participants to the submarket $\Omega(y)$ are randomly placed on the circumference of a circle according to a uniform distribution.
- Each jobseeker applies clockwise to the nearest firm.
- The probability that the firm receives n applications follows a geometric distribution: $\frac{1}{1+\lambda}\left(\frac{\lambda}{1+\lambda}\right)^n$. The mean of the geometric distribution is exactly λ.

Modelling the search location friction

- Microfoundation: Participants to the submarket $\Omega(y)$ are randomly placed on the circumference of a circle according to a uniform distribution.
- Each jobseeker applies clockwise to the nearest firm.
- The probability that the firm receives n applications follows a geometric distribution: $\frac{1}{1+\lambda}\left(\frac{\lambda}{1+\lambda}\right)^n$. The mean of the geometric distribution is exactly λ.

- Add a parameter κ to vary the intensity of the **location friction**.
- Think of κ as the probability that at the end of the screening process the candidate is still interested and accepts the offer
- Firms receive, on average, $\kappa\lambda(y)$ employable applicants.

Modelling the screening friction

- The applicants to firm y are randomly ordered in a queue.
- The firm interviews the first applicant and learns her type.
- With probability $\sigma \in [0, 1]$ the firm can interview another applicant and learn her type.

- σ is the **screening friction**: firms may fail to identify the high quality of some applicants.

Bringing the two together

- The probability that a firm y in submarket $\Omega(y)$ screens *at least one* worker of high type - and this worker accepts the offer - can be written as:

$$\phi(\mu, \lambda) = \kappa \frac{\mu}{1 + \sigma\mu + (1 - \sigma)\lambda} \tag{1}$$

- The probability that the firm screens at least one employable worker that accepts the offer is: $\phi(\lambda, \lambda) = m(\lambda) = \kappa\lambda/(1 + \lambda)$

Hiring

- The firm hires the screened applicant with the highest type.
- A match between firm y and worker x produces net output $f(x, y)$.

Hiring

- The firm hires the screened applicant with the highest type.
- A match between firm y and worker x produces net output $f(x, y)$.

- The expected surplus is therefore:

$$S(\mu, \lambda, y) = m(\lambda) f(x_L, y) + \phi(\mu, \lambda) \left[f(x_H, y) - f(x_L, y) \right] \tag{2}$$

Outline

1. What constraints labor market matching in LICs?

2. **How can we raise total employment and match quality?**
   2.1 Certification and tests
   2.2 Elicitation of preferences
   2.3 Recommendation algorithm

3. How can we identify total employment and match-quality effects?

Raising κ and σ will increase employment and match quality

$$\phi(\mu, \lambda) = \kappa \frac{\mu}{1 + \sigma\mu + (1-\sigma)\lambda} \tag{3}$$

$$S(\mu, \lambda, y) = m(\lambda) f(x_L, y) + \phi(\mu, \lambda) \left[f(x_H, y) - f(x_L, y) \right] \tag{4}$$

Our experiment

- We run a two-sided experiment with firms and jobseekers in Addis Ababa, Ethiopia

Our experiment

- We run a two-sided experiment with firms and jobseekers in Addis Ababa, Ethiopia
- We sample the universe of firms posting vacancies in 4 white-collar occupations

Our experiment

- We run a two-sided experiment with firms and jobseekers in Addis Ababa, Ethiopia

- We sample the universe of firms posting vacancies in 4 white-collar occupations
- We recruit active jobseekers interested in jobs in these occupations

Our experiment

- We run a two-sided experiment with firms and jobseekers in Addis Ababa, Ethiopia

- We sample the universe of firms posting vacancies in 4 white-collar occupations
- We recruit active jobseekers interested in jobs in these occupations

- We measure:
  - Firms' preferences over worker attributes and jobseekers' preferences over job attributes.
  - Jobseeker skills (in great detail)
  - Outcomes three months after treatment

Two treatment arms: Recommendations & Certificate

- In T1 we aim to raise κ: we share with the jobseeker a personalised list of vacancies that are a good match for her profile
  - We decrease the location friction by encouraging jobseekers to apply to more of their close-by vacancies.
- In T2 we aim to raise σ: we offer T1 & provide the jobseeker with a skill certificate
  - Probability firm screens one more applicant $\uparrow$
- We cross-randomize the two treatments (ie, we split each group in C, T1, T2 and only make recommendations within group).

Outline

1. What constraints labor market matching in LICs?

2. **How can we raise total employment and match quality?**
   2.1 Certification and tests
   2.2 Elicitation of preferences
   2.3 Recommendation algorithm

3. How can we identify total employment and match-quality effects?

Tests

Each jobseeker completes a battery of **cognitive**, **personality**, and **skill** tests.

Domain	**Test**	**Targeted construct**
Cognitive	Raven's Progressive Matrices	General intelligence
Hard Skill	Mathematics	Mathematics proficiency
Hard Skill	Amharic literacy	Amharic language proficiency
Hard Skill	English	English language proficiency
Personality	Big Five (BFI–44)	Conscientiousness
Personality	Big Five (BFI–44)	Emotional control
Personality	Grit (12 items)	Perseverance
Social	Reading the Mind in the Eyes	Perspective taking
Cognitive	Stroop	Executive function
Personality	Self-efficacy	Confidence / initiative
Personality	Growth mindset	Learning orientation

Certificate

JOBSEEKER SKILLS CERTIFICATE

Issued to

[redacted]

Top-third in 7/11 skills

This certificate helps jobseekers present their skills and find jobs that fit them well. It is provided by TalentConnect, a partnership between researchers at the University of Oxford, Peking University, and Vrije Universiteit Amsterdam. See the second page for detailed explanations and how performance bands are defined.

HARD SKILLS	Performance band
General intelligence: Shows strong reasoning and problem-solving ability.	Top third ●●●
Mathematics proficiency: Can use practical mathematics and numerical reasoning in work tasks.	Top third ●●●
Amharic proficiency: Can understand and communicate clearly in Amharic for workplace tasks.	Top third ●●●
English proficiency: Can understand and use English for workplace communication.	Top third ●●●
SOFT SKILLS	
Conscientiousness: Is diligent and can be relied upon to follow instructions.	Bottom third ●○○
Emotional control: Is emotionally stable and resilient.	Middle third ●●○
Grit / perseverance: Keeps going and does not give up when work is difficult.	Top third ●●●
Perspective taking: Understands others and works smoothly with colleagues or clients.	Top third ●●●
Executive function: Stays focused and avoids mistakes under time pressure.	Top third ●●●
Self-efficacy: Is confident and takes initiative without close supervision.	Middle third ●●○
Growth mindset: Is eager to learn and improve when facing challenges.	Bottom third ●○○

Issue date: 06 May 2026

L. Hensel

Professor Lukas Hensel
(Principal Investigator, TalentConnect)

Verification:
+251972239787 | +251972154787

Click here to verify

Outline

1. What constraints labor market matching in LICs?

2. How can we raise total employment and match quality?

   2.1 Certification and tests

   2.2 Elicitation of preferences

   2.3 Recommendation algorithm

3. How can we identify total employment and match-quality effects?

Discrete choice experiment

- Identify manager in the firm that has most control on the hiring process.
- Each manager completes 10 binary choice tasks in a discrete choice experiment (DCE)
- With multiple vacancies, manager completes a separate DCE for each vacancy
- We instruct respondents to assume that all candidates are identical except for the characteristics explicitly mentioned in the exercise

- We run a similar DCE for jobseekers Jobseeker DCE.

What information do we elicit with the DCE?

- In each task, the hiring manager chooses between two hypothetical applicants that differ along 8 randomly varied attributes:
  - Education, experience, location, gender, 1st soft skill, 2nd soft skill, random soft skill, hard skill

- After selecting the candidate, we ask the respondent to also state:
  - The maximum monthly wage she would be willing to pay for each candidate
  - The likelihood the candidate will accept the job
  - Willingness to pay to receive one extra candidate with those characteristics

Occupation specific preferences

- We estimate vacancy v's preference vector β_v as:

$$\beta_v = \underbrace{b^F}_{\text{population mean}} + \underbrace{\gamma_{o(v)}}_{\text{occupation shift}} + \underbrace{\eta_v}_{\text{vacancy-specific deviation}}, \qquad \eta_v \sim \mathcal{N}(0, \text{diag}(\sigma^F))$$

- "Vacancy" utility can then be written as:

$$U_{v,j} = \beta_v W_j + \varepsilon_{v,j}$$

with $\varepsilon \sim$ i.i.d. and W_j a vector of jobseeker j's characteristics

- Two-step process:
  1. Step 1: $(b^F, \{\gamma_o\}_{o \neq \text{ref}}, \sigma^F)$ are estimated by simulated maximum likelihood
  2. Step 2: vacancy specific deviations η_v are estimated by empirical Bayes Details

Estimated coefficients

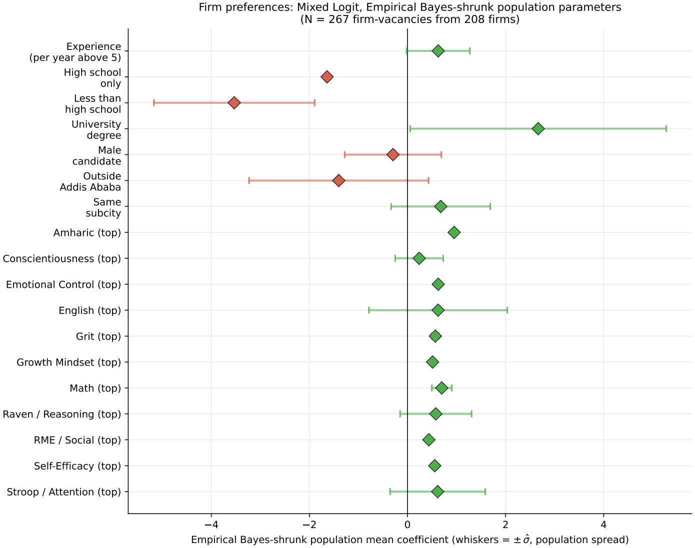

Outline

The recommendation algorithm

- The sample is partitioned into independent sub-markets by treatment group $\{C, T1, T2\}$ and target occupation.
- We use a deferred acceptance algorithm to identify the recommendations. Details
  - Start from the cardinal preferences identified in the DCE.
  - Convert them in ordinal rankings.
  - Choose jobseeker-proposing or firm-proposing based on total ‘utility produced’
- Jobseekers receive the recommendations by SMS and are then nudged to apply with a phone call.

Are recommended candidates better than non-recommended?

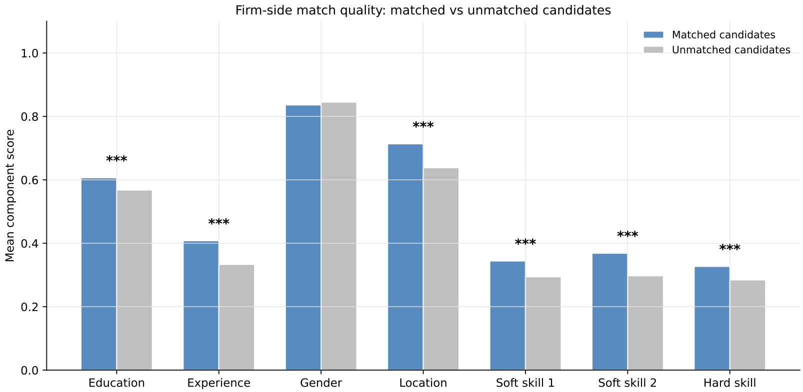

Result: Recommended candidates dominate *every* component of firm utility (except gender). Selection statistics Recommendation statistics

Outline

1. What constraints labor market matching in LICs?

2. How can we raise total employment and match quality?
   - 2.1 Certification and tests
   - 2.2 Elicitation of preferences
   - 2.3 Recommendation algorithm

3. How can we identify total employment and match-quality effects?

Worker-level, vacancy-level, and market-level experiments

We cannot identify total employment effects by running

- worker-level experiments
  - Untreated workers may be displaced by the intervention.
- vacancy-level experiments
  - Untreated vacancies may not be filled due to the intervention.

Market-level (e.g., city level) experiments would identify equilibrium effects, but

- These experiments are very costly to implement.
- And are near impossible in low-income countries with weak employment agencies.

Two-sided experiments, dyadic data, structural estimation

We can identify a *lower bound* to total employment effects if we:

- Introduce two-sided randomization
- Collect dyadic data
- Assume displacement only occurs as a result of a unit matching with another unit (e.g., applications that do not result in a match do not cause displacement).
- Assume there is no displacement in the control group.

To go beyond the lower bound and identify the *exact magnitude* of the impact on total employment, we will need a structural model.

A simple illustration

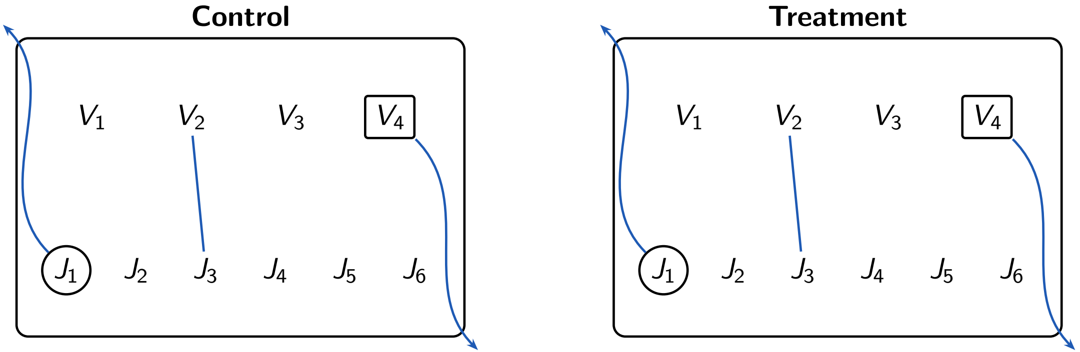

- **Blue** edges: matches that would have formed regardless of treatment

Unobserved worker displacement

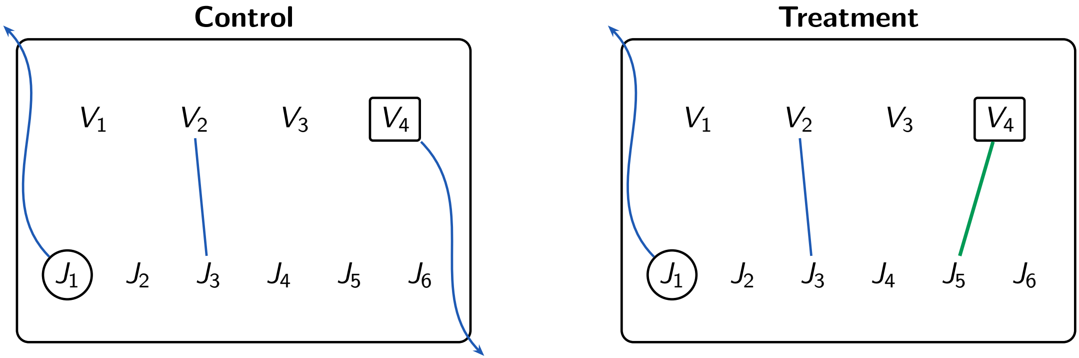

- We add the **Green** edge → **observed employment** ↑, but **worker displacement**
- With two-sided randomization and dyadic data we can identify direct displacement
- ... an upper bound to total displacement as the displaced worker may rematch.
- In this case, the lower bound to impact on total matches is 0.

Unobserved vacancy displacement

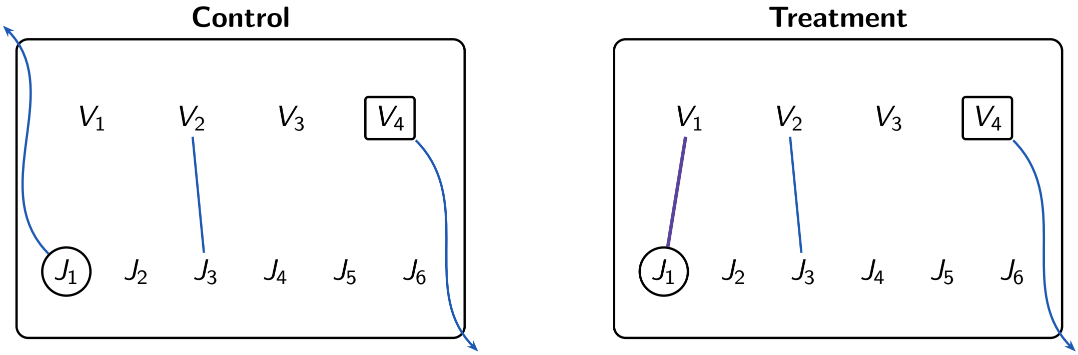

- We add the **Violet** edge $\rightarrow$ **observed vacancy filling** $\uparrow$, but **vacancy displacement**
- Again we can bound total displacement effect and impact on total matches

A truly additional match

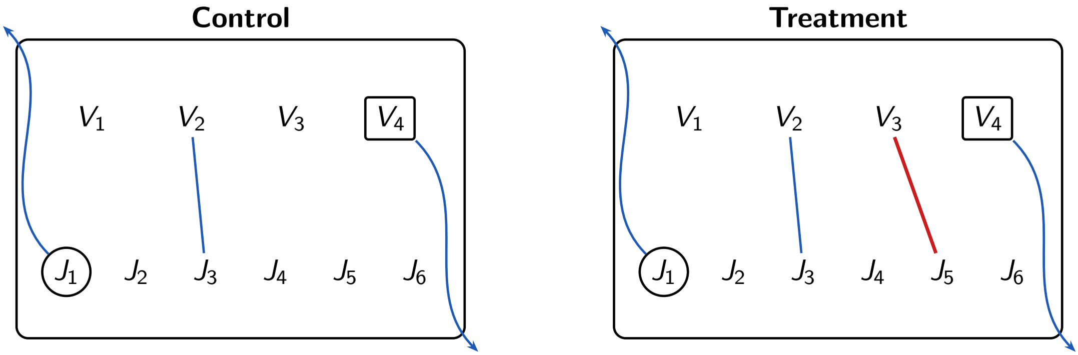

- We add the **Red** edge → **observe vacancy filling, employment** ↑; **no displacement**
- Here we can correctly conclude we have increased total matches.

Two-sided rand Match quality

Control displacement - absent intervention

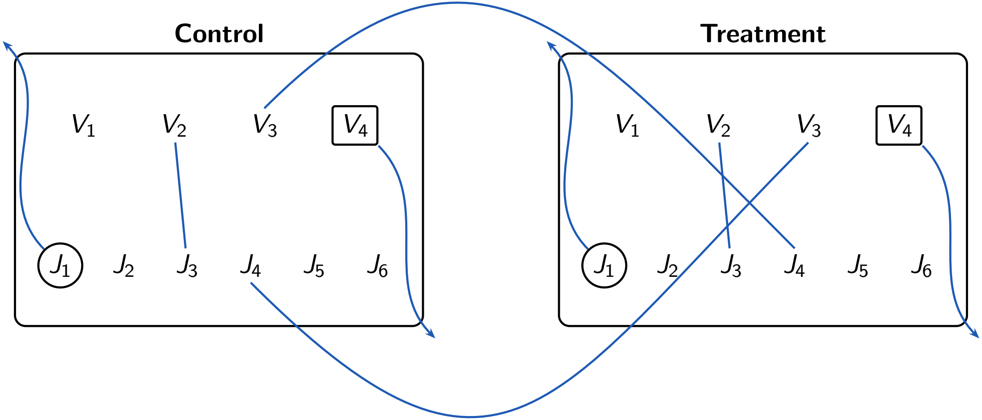

Control displacement - with intervention

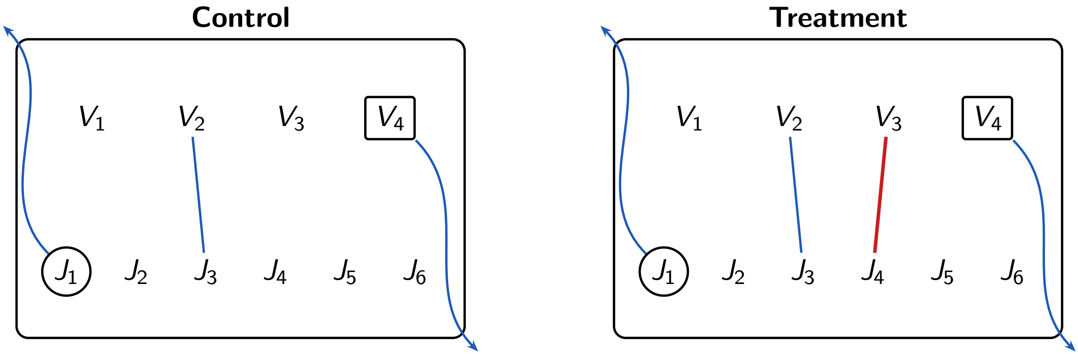

- Observed **vacancy filling, employment** $\uparrow$; **control displacement**
- We introduce super-control weeks to test for control displacement.

Back to the structural model: calibration of auxiliary parameters

$$\phi(\mu, \lambda) = \kappa \frac{\mu}{1 + \sigma\mu + (1 - \sigma)\lambda}$$

$$S(\mu, \lambda, y) = m(\lambda) f(x_L, y) + \phi(\mu, \lambda) \left[f(x_H, y) - f(x_L, y) \right]$$

Back to the structural model: calibration of auxiliary parameters

$$\phi(\mu, \lambda) = \kappa \frac{\mu}{1 + \sigma\mu + (1 - \sigma)\lambda}$$

$$S(\mu, \lambda, y) = m(\lambda) f(x_L, y) + \phi(\mu, \lambda) \left[f(x_H, y) - f(x_L, y) \right]$$

- We can calibrate μ and λ from population statistics, properly weighted.
- $f(x, y)$ is given by the WTP question in the DCE

Back to the structural model: calibration of the frictions and the shocks to the frictions

$$\phi(\mu, \lambda) = \kappa \frac{\mu}{1 + \sigma\mu + (1-\sigma)\lambda}$$

$$S(\mu, \lambda, y) = m(\lambda) f(x_L, y) + \phi(\mu, \lambda) \left[f(x_H, y) - f(x_L, y) \right]$$

- κ is identified by the vacancy filling rate in the control group;
- σ is identified by the high-quality vacancy filling rate in the control group;
- κ' and σ' are identified by the difference in vacancy filling rates and high-quality filling rates in T1 vs C, and T2 vs C.

Bibliography I

Balgova, Maria et al. (2025). *Wage Information and Applicant Selection*.

Cai, Xiaoming, Pieter Gautier, and Ronald Wolthoff (2023). "Meetings and Mechanisms". In: *International Economic Review* 64.1, pp. 155–185. DOI: 10.1111/iere.12592.

— (2025). "Search, Screening, and Sorting". In: *American Economic Journal: Macroeconomics* 17.3, pp. 205–236. DOI: 10.1257/mac.20240026.

Firm requests a considerable number of additional applicants

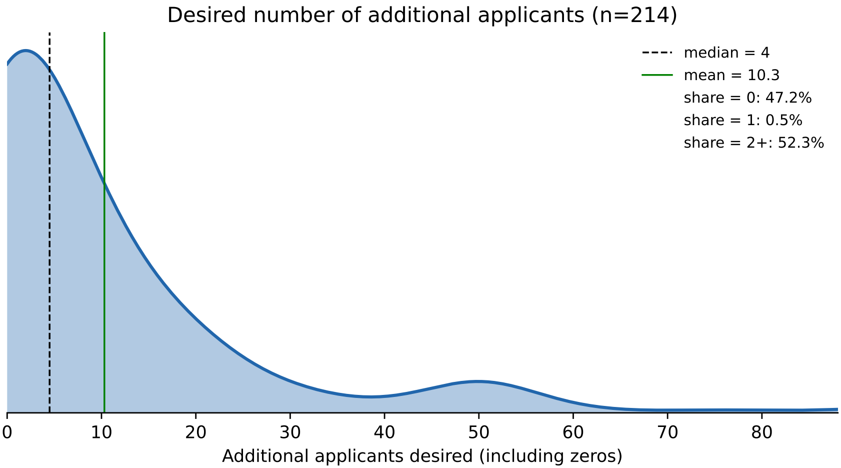

Back

Discrete choice experiment - Jobseekers Back

- Each jobseeker completes 10 binary choice tasks in a discrete choice experiment (DCE)
- In each task, the jobseeker chooses between 2 hypothetical vacancies that differ along 7 randomly varied attributes, selecting the one she would apply to:
  - Location: *Central / Non Central*
  - Job type: *Permanent / Temporary*
  - Weekly working hours: *30 / 40 / 50*
  - Paid leave days: *0 / 10 / 20*
  - Training offered: *Yes / No*
  - Additional benefits (like health insurance): *Yes / No*
  - Schedule flexibility: *Fixed / Flexible*

Discrete choice experiment - Jobseekers Back

- Before the exercise, we instruct respondents to assume that all attributes not shown (e.g., the offered wage) are held constant across the two options
- After selecting the preferred vacancy, we ask the respondent to also state:
  - What's the minimum monthly wage she would accept for each of these jobs (in ETB gross)
  - What's the likelihood of getting the job (conditional on having applied)
  - How long do they expect to remain employed in this job (expected tenure conditional on getting hired)

Estimation strategy - Jobseekers Back

- We want to recover an individual preference vector $\hat{\beta}_j$ for each jobseeker j over the stated vacancy attributes

Estimation strategy - Jobseekers Back

- We want to recover an individual preference vector $\hat{\beta}_j$ for each jobseeker j over the stated vacancy attributes
- We exclude wages because pay is frequently not disclosed in job ads. The choice tasks mirror the information set typically available at the application stage (Balgova et al., 2025)

Estimation strategy - Jobseekers Back

- We want to recover an individual preference vector $\hat{\beta}_j$ for each jobseeker j over the stated vacancy attributes
- We exclude wages because pay is frequently not disclosed in job ads. The choice tasks mirror the information set typically available at the application stage (Balgova et al., 2025)
- We use the individual-specific estimated preference vector $\hat{\beta}_j$ to rank vacancies. These rankings feed into the matching algorithm that generates vacancy recommendations

Estimation strategy - Jobseekers Back

- We want to recover an individual preference vector $\hat{\beta}_j$ for each jobseeker j over the stated vacancy attributes
- We exclude wages because pay is frequently not disclosed in job ads. The choice tasks mirror the information set typically available at the application stage (Balgova et al., 2025)
- We use the individual-specific estimated preference vector $\hat{\beta}_j$ to rank vacancies. These rankings feed into the matching algorithm that generates vacancy recommendations
- From the DCE we observe a sequence of choices for each jobseeker j, y_j

Estimation strategy - Jobseekers Back

- We estimate a mixed logit (random-coefficients logit) on the binary choice data

Estimation strategy - Jobseekers Back

- We estimate a mixed logit (random-coefficients logit) on the binary choice data
- Individual-specific preferences satisfy $\beta_j \sim \mathcal{N}(b^{JS}, \Sigma^{JS})$

Estimation strategy - Jobseekers Back

- We estimate a mixed logit (random-coefficients logit) on the binary choice data
- Individual-specific preferences satisfy $\beta_j \sim \mathcal{N}(b^{JS}, \Sigma^{JS})$
- **Step 1 (population)**: estimate hyperparameters (b^{JS}, Σ^{JS}) by simulated maximum likelihood

Estimation strategy - Jobseekers Back

- We estimate a mixed logit (random-coefficients logit) on the binary choice data
- Individual-specific preferences satisfy $\beta_j \sim \mathcal{N}(b^{JS}, \Sigma^{JS})$
- **Step 1 (population)**: estimate hyperparameters (b^{JS}, Σ^{JS}) by simulated maximum likelihood
- **Step 2 (individual)**: compute empirical Bayes posterior means

$$\hat{\beta}_j = \mathbb{E}[\beta_j \mid y_j, \hat{b}^{JS}, \hat{\Sigma}^{JS}],$$

i.e., the posterior under plug-in hyperparameters, evaluated numerically

Estimation strategy - Jobseekers Back

- We estimate a mixed logit (random-coefficients logit) on the binary choice data
- Individual-specific preferences satisfy $\beta_j \sim \mathcal{N}(b^{JS}, \Sigma^{JS})$
- **Step 1 (population)**: estimate hyperparameters (b^{JS}, Σ^{JS}) by simulated maximum likelihood
- **Step 2 (individual)**: compute empirical Bayes posterior means

$$\hat{\beta}_j = \mathbb{E}[\beta_j \mid y_j, \hat{b}^{JS}, \hat{\Sigma}^{JS}],$$

  i.e., the posterior under plug-in hyperparameters, evaluated numerically
- Note that in practice we restrict the random-coefficient variance–covariance matrix to be diagonal ($\Sigma^{JS} = diag(\sigma^{JS})$), which rules out correlation in tastes across vacancy attributes. Given only 10 choice tasks per respondent, estimating a full matrix would be weakly identified and likely yield unstable estimates

Estimated coefficients

Back

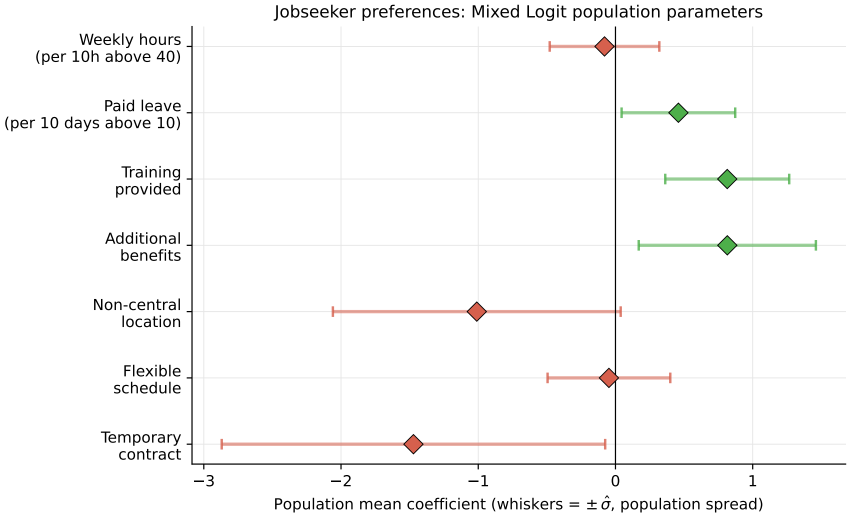

Selection statistics Back

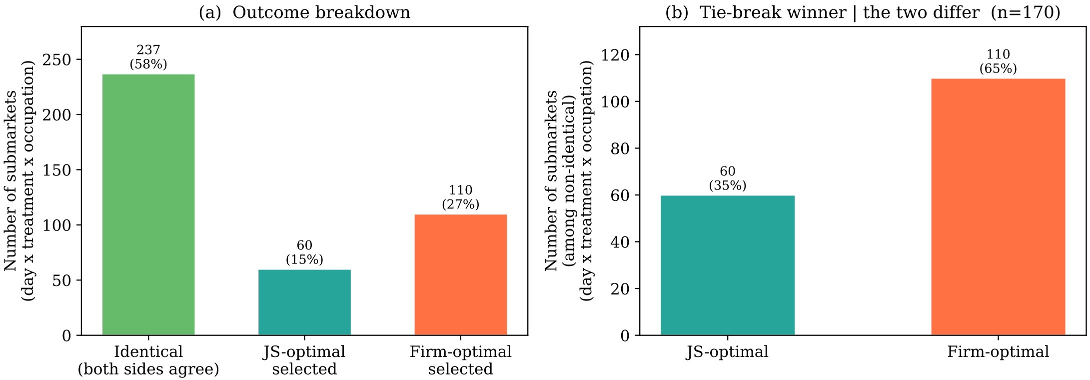

Selection statistics Back

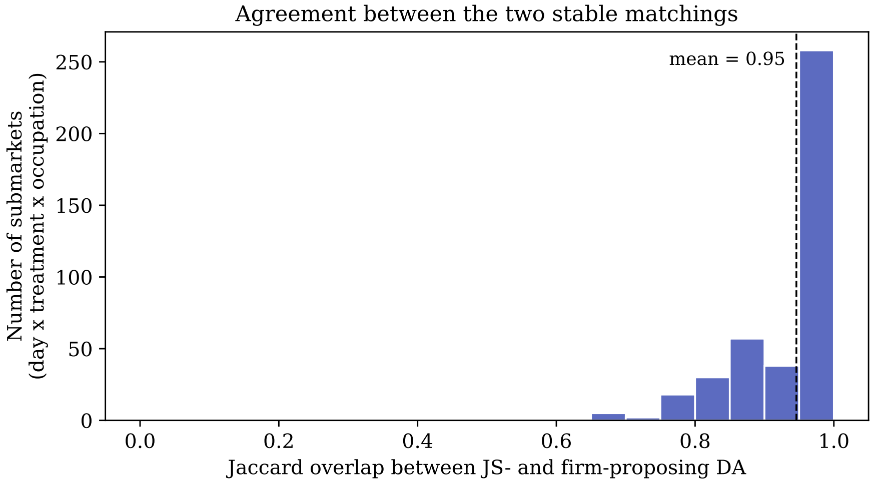

Recommendation statistics

Back

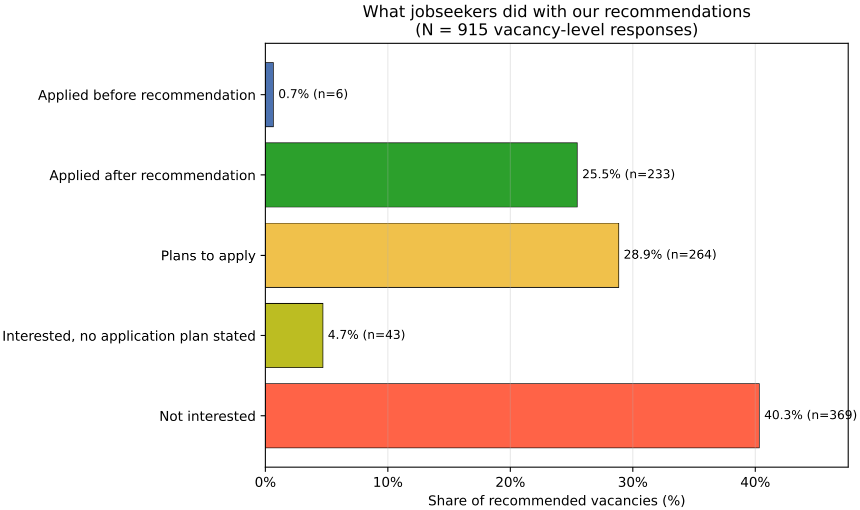

Recommendation statistics Back

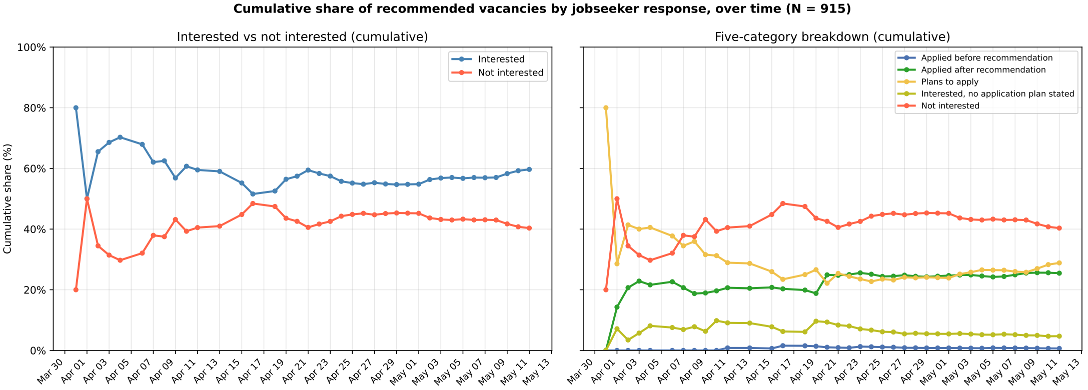

Recommendation statistics Back

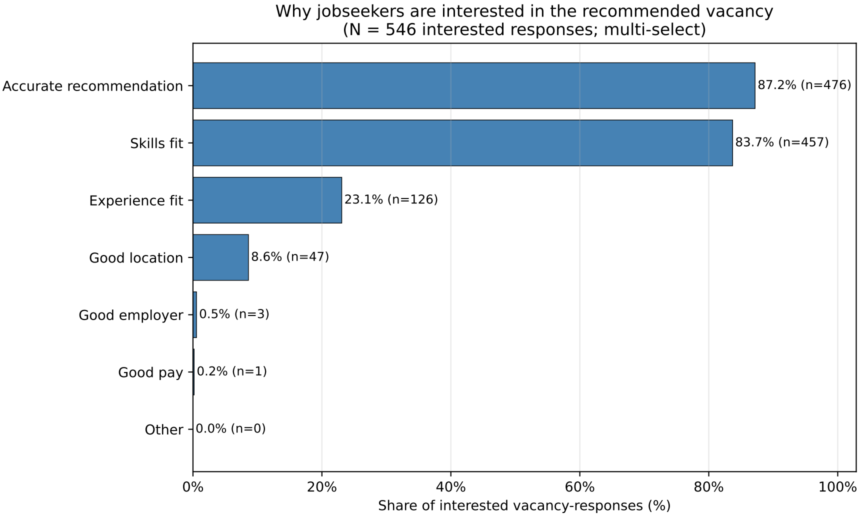

We need both two-sided randomization and dyadic data [Back]

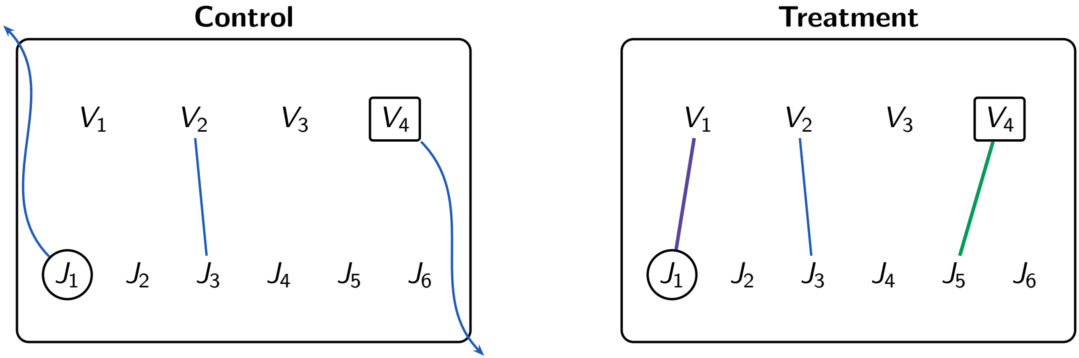

- **Observed vacancy filling, employment** ↑, but vacancy and worker displacement
- Observing two-sided impacts insufficient to conclude total matches ↑
- True impact on total matches could be zero

Improvement in match quality Back

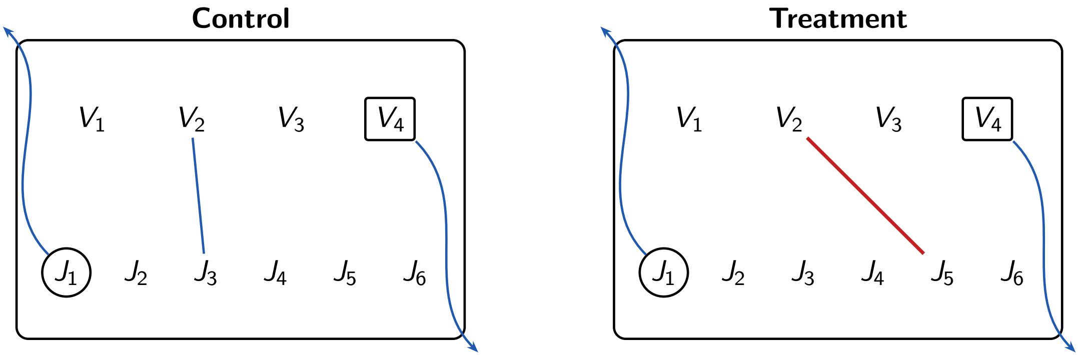

- No change in vacancy filling, employment
- Better match for V_2 (assuming J_3's application is not crowded out)
- Can identify this using measures of match quality

The recommendation market Back

- The market consists of:
  - A set $\mathcal{J}$ of jobseekers, each with preferences ordering over vacancies
  - A set $\mathcal{V}$ of vacancies, each with preferences ordering over jobseekers
  - **Quotas**: each vacancy v can receive up to q_v candidates (8 or 16 depending on the number of open positions); each jobseeker receives up to $q_j = 5$ recommendations.

The recommendation market Back

- The market consists of:
  - A set $\mathcal{J}$ of jobseekers, each with preferences ordering over vacancies
  - A set $\mathcal{V}$ of vacancies, each with preferences ordering over jobseekers
  - **Quotas**: each vacancy v can receive up to q_v candidates (8 or 16 depending on the number of open positions); each jobseeker receives up to $q_j = 5$ recommendations.
- The market is partitioned into independent sub-markets by treatment group $\{C, T1, T2\}$ and target occupation.

The recommendation market Back

- The market consists of:
  - A set $\mathcal{J}$ of jobseekers, each with preferences ordering over vacancies
  - A set $\mathcal{V}$ of vacancies, each with preferences ordering over jobseekers
  - **Quotas**: each vacancy v can receive up to q_v candidates (8 or 16 depending on the number of open positions); each jobseeker receives up to $q_j = 5$ recommendations.
- The market is partitioned into independent sub-markets by treatment group $\{C, T1, T2\}$ and target occupation.
- The recommendation algorithm runs separately within each sub-market

Ordinal rankings within an occupation Back

- For each jobseeker j we estimate preferences $\hat{\beta}_j$ over job attributes X_v, and for each vacancy v we estimate hiring-side preferences $\hat{\beta}_v$ over worker attributes W_j (from the vacancy-specific DCE)

Ordinal rankings within an occupation Back

- For each jobseeker j we estimate preferences $\hat{\beta}_j$ over job attributes X_v, and for each vacancy v we estimate hiring-side preferences $\hat{\beta}_v$ over worker attributes W_j (from the vacancy-specific DCE)
- We then compute match-specific *systematic utilities* (scores), omitting the idiosyncratic error term:

$$\tilde{U}_j(v) = \hat{\beta}_j X_v \quad \forall v \in \mathcal{V}_o \tag{5}$$

$$\tilde{U}_v(j) = \hat{\beta}_v W_j \quad \forall j \in \mathcal{J}_o \tag{6}$$

where $\mathcal{V}_o$ and $\mathcal{J}_o$ denote the set of n vacancies and m jobseekers respectively in occupation (submarket) o (and treatment group t)

Ordinal rankings within an occupation Back

- For each jobseeker j we estimate preferences $\hat{\beta}_j$ over job attributes X_v, and for each vacancy v we estimate hiring-side preferences $\hat{\beta}_v$ over worker attributes W_j (from the vacancy-specific DCE)
- We then compute match-specific *systematic utilities* (scores), omitting the idiosyncratic error term:

$$\tilde{U}_j(v) = \hat{\beta}_j X_v \quad \forall v \in \mathcal{V}_o \tag{5}$$

$$\tilde{U}_v(j) = \hat{\beta}_v W_j \quad \forall j \in \mathcal{J}_o \tag{6}$$

  where $\mathcal{V}_o$ and $\mathcal{J}_o$ denote the set of n vacancies and m jobseekers respectively in occupation (submarket) o (and treatment group t)
- Using these scores, we can construct ordinal rankings for each side of the market:

$$\tilde{U}_j(v_{(1)}) \geq \tilde{U}_j(v_{(2)}) \geq \cdots \geq \tilde{U}_j(v_{(n)}) \quad \forall j \in \mathcal{J}_o \tag{7}$$

$$\tilde{U}_v(j_{(1)}) \geq \tilde{U}_v(j_{(2)}) \geq \cdots \geq \tilde{U}_v(j_{(m)}) \quad \forall v \in \mathcal{V}_o \tag{8}$$

Why use a two-sided Deferred Acceptance (DA) algorithm?

Back

Definition (Stable matching)

A matching μ is *stable* if:

(i) **Individual rationality**: every matched pair $(j, v) \in \mu$ is acceptable to both sides.

(ii) **No blocking pair**: there is no $(j, v) \notin \mu$ such that j prefers v to at least one current match, or has spare capacity, *and* v prefers j to at least one current candidate, or has spare capacity.

- **Commensurability:** $\tilde{U}_j(\cdot)$ and $\tilde{U}_v(\cdot)$ are estimated from different models and therefore live on different scales.
- **Stability:** we want to avoid recommendations that generate *blocking pairs*, i.e. jobseeker–vacancy pairs that would mutually prefer each other to their assigned match.
- DA only requires *ordinal* rankings on each side of the market, making the procedure scale-free and producing stable match recommendations.

Selecting between Jobseeker and Vacancy-Optimal Matchings

Back

JS-proposing DA

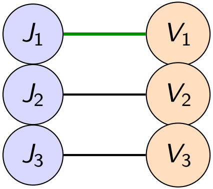

$$\mathrm{JointUtility}_{\mathsf{JS}} = +1$$

Firm-proposing DA

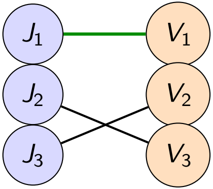

$$\mathrm{JointUtility}_{\mathsf{Firm}} = +0.5$$

Joint utility — sum over *all* matched pairs of z-standardised scores (within sub-market) so that firm & jobseeker scales are comparable:

$$\mathrm{JU}(\mathcal{M}) = \sum_{(j,v)\in\mathcal{M}} \left[\tilde{U}_j(v) + \tilde{U}_v(j)\right]$$

$J_1 - V_1$ appears in *both* matchings (*fixed point*) → only differing pairs tilt the comparison.

Rule: if $\mathrm{JU}_{\mathsf{JS}} > \mathrm{JU}_{\mathsf{Firm}} \Rightarrow$ keep **JS-optimal allocation**, otherwise keep **Firm-optimal allocation** (ties → JS)

Firm-level clustering and empirical Bayes shrinkage Back

- The DCE is observed at the vacancy level, but vacancies within the same firm share unobserved firm-level heterogeneity. We cluster standard errors of the SML estimator $(\hat{b}^F, \{\hat{\gamma}_o\}, \hat{\sigma}^F)$ at the firm level $f(v)$.

Firm-level clustering and empirical Bayes shrinkage Back

- The DCE is observed at the vacancy level, but vacancies within the same firm share unobserved firm-level heterogeneity. We cluster standard errors of the SML estimator $(\hat{b}^F, \{\hat{\gamma}_o\}, \hat{\sigma}^F)$ at the firm level $f(v)$.
- **Step 2 (individual posteriors).** For each vacancy v, the vacancy-specific deviation is recovered as the posterior mean under plug-in hyperparameters:

$$\hat{\eta}_v \;=\; \mathbb{E}\big[\eta_v \,\big|\, y_v,\ \hat{b}^F,\ \{\hat{\gamma}_o\},\ \hat{\sigma}^F\big], \qquad \hat{\beta}_v \;=\; \hat{b}^F + \hat{\gamma}_{o(v)} + \hat{\eta}_v.$$

Firm-level clustering and empirical Bayes shrinkage Back

- The DCE is observed at the vacancy level, but vacancies within the same firm share unobserved firm-level heterogeneity. We cluster standard errors of the SML estimator $(\hat{b}^F, \{\hat{\gamma}_o\}, \hat{\sigma}^F)$ at the firm level $f(v)$.
- **Step 2 (individual posteriors).** For each vacancy v, the vacancy-specific deviation is recovered as the posterior mean under plug-in hyperparameters:

$$\hat{\eta}_v \;=\; \mathbb{E}\big[\eta_v \,\big|\, y_v,\; \hat{b}^F,\; \{\hat{\gamma}_o\},\; \hat{\sigma}^F\big], \qquad \hat{\beta}_v \;=\; \hat{b}^F + \hat{\gamma}_{o(v)} + \hat{\eta}_v.$$

- **Empirical Bayes shrinkage.** The posterior mean of $\eta_{v,k}$ is a precision-weighted average of the data-implied estimate and the prior mean (zero):

$$\hat{\eta}_{v,k} \;=\; \omega_{v,k}\,\tilde{\eta}_{v,k}, \qquad \omega_{v,k} = \frac{\tau_{v,k}}{\tau_{v,k} + 1/\hat{\sigma}_k^2}.$$

  When the DCE choices are uninformative about attribute k (low $\tau_{v,k}$), $\omega_{v,k} \to 0$ and $\hat{\beta}_{v,k}$ is pulled toward the occupation-adjusted population mean $\hat{b}_k^F + \hat{\gamma}_{o(v),k}$