

Recruiting, Screening, and Entry into Organizations

Organizations, Personnel, and Internal Labor Markets

Week 2

Sai Zhang

Hiring begins before the firm chooses a worker

Central question

How do firms choose whom to evaluate and hire when attention, time, information, and interview capacity are costly?

- Equal application counts can conceal different review rules, shortlists, offers, and acceptance decisions.
- Recruiting changes both the opportunity set and the information used to choose within it.
- Follow employer decisions and applicant responses together.

Central case: Hoffman–Kahn–Li (2018), *Discretion in Hiring*. Field map: Hoffman–Stanton (2025), published handbook chapter.

Learning goals: defend a recruiting comparison

- Define applications, assessments, offers, and starts with distinct risk sets and horizons.
- Calculate the value of another evaluation under explicit assumptions.
- Compare evidence on tests, discretion, order, referrals, and recruiting effort.
- Explain what missing performance for rejected applicants prevents us from learning.

Core points

Recruiting is sequential; the observed population changes at each stage. Faster filling, average hire quality, and worker welfare are different objectives. A credible design names the comparison and the missing outcomes.

Keep application, offer, and start as separate events

Outreach $\longrightarrow$ application $\longrightarrow$ assessment
 $\longrightarrow$ offer $\longrightarrow$ acceptance $\longrightarrow$ start

- A person may submit several applications; a vacancy may contain several positions.
- Retain applicant, application, vacancy, position, and event identifiers.
- Direct offers or multiple assessment routes require branching states.
- Never equate a pending case with rejection, an unknown outcome with zero, or an accepted offer with a start.

A stage ledger makes the denominator visible

Event	Useful denominator	Record that preserves the claim
Application	Exposed/eligible audience, if known	Channel and exposure date
Assessment	Applications eligible for review	Queue, invitation, no-show, score
Offer	Baseline applications; assessed cases	Terms, decision and offer times
Acceptance	Offers made in the cohort	Response, withdrawal, deadline
Start	Accepted offers; baseline applications	Start date and cancelled positions
Retention	Entrants at a stated horizon	Exit time and administrative censoring

State cohort entry, follow-up end, competing exits, and missingness for every row.

The funnel product is an accounting identity

$$\frac{N_K}{N_0} = \prod_{s=1}^K \frac{N_s}{N_{s-1}}.$$

- N_0 : original applications; N_s : cases reaching stage s ; K : final stage.
- Events must be nested within the same cohort and compatible follow-up horizons; denominators must be defined and nonzero.
- The intermediate counts cancel. No causal assumption is supplied by that algebra.

Counterexample to diagnose

This month's screening rate $\times$ last month's offer acceptance $\times$ annual retention does not describe one cohort's probability of a retained start.

A higher stage conversion rate has several explanations

- Interview-to-offer conversion can rise because interviews improve information, the shortlist strengthens, or the offer threshold falls.
- Reaching the interview is potentially changed by the policy itself.
- A conditional transition describes the selected stage population; it is not automatically a causal mechanism.

Preserve the original risk set

Report starts per baseline application alongside conditional transitions. A randomized recruiting advertisement may also change who applies, so its audience can be a broader relevant denominator.

Another evaluation is valuable only if it changes the choice

- **Teaching benchmark:** one opening, recall of evaluated candidates, finite horizon, and known exogenous candidate-value draws.
- b : best known net match surplus, including the option to leave the opening unfilled.
- Evaluating the next candidate costs c_t and reveals value Q_{t+1} .
- A noisy test is information about value. A ranking needs a justified cardinal mapping before entering a surplus calculation.

Conceptual ancestry: Weitzman (1979). This simplified stopping model is a pedagogical construction, not an estimated hiring model.

The stopping rule compares acting now with learning more

$$V_t(b) = \max \{b, -c_t + \delta \mathbb{E}[V_{t+1} \{\max(b, Q_{t+1})\}]\}, \quad V_T(b) = b.$$

- $t = 0, \dots, T - 1$; V_t : value of remaining opportunities; T : terminal date; $\delta \in (0, 1]$: discount factor. Assume b is a sufficient state.
- Stopping takes b . Continuing pays the cost and preserves recall of the best candidate.
- At T , there is no further evaluation: take the best available option.
- Competing employers, withdrawal, strategic acceptance, uncertain arrival, and multiple slots require additional states or assumptions.

At cost 1, continue; at cost 3, stop

$$b = 8, \quad \delta = 1, \quad Q = \begin{cases} 6 & \text{with probability } 1/2, \\ 12 & \text{with probability } 1/2. \end{cases}$$

$$\mathbb{E}[\max(8, Q)] = \frac{1}{2} 8 + \frac{1}{2} 12 = 10.$$

Evaluation cost	Stop now	Evaluate once	Choice
1	8	$10 - 1 = 9$	Continue
3	8	$10 - 3 = 7$	Stop

Early closure can reflect a strong current candidate or costly delay; it does not alone reveal low standards.

All values and probabilities are hypothetical teaching inputs. One evaluation remains.

An attractive candidate must still accept the offer

$$S_i(w, d) = a_i(w, d) \bar{v}_i(w, d) - k_i(w, d).$$

- w : compensation; d : delay; a_i : acceptance probability.
- $\bar{v}_i$: expected employer surplus *conditional on acceptance*; k_i : offer or delay cost.
- The highest test score need not yield the largest expected realized hiring surplus.
- This notation identifies no acceptance elasticity. Worker welfare requires separate pay, job-condition, and participation-cost outcomes.

Pedagogical offer calculation. Acceptance conditioning is part of the definition of $\bar{v}_i$.

Recruiting channels determine whom the firm can consider

- **Setting:** Weinstein studies on-campus recruiting with employer screening costs in an equilibrium framework.
- **Economic choice:** where to recruit determines the cost of reaching and evaluating potential workers.
- **Objects:** channels considered, campus visits, applications, interview costs, and hires.
- **Limit:** a regression only among visited campuses conditions on the recruiting choice being explained; model counterfactuals retain their assumptions.

Data request

Record channels not chosen and eligible populations where known, alongside the channels that produced hires.

Weinstein (2018), *Employer Screening Costs, Recruiting Strategies, and Labor Market Outcomes*.

Pay information can attract applicants and signal the job

	Dal Bó–Finan–Rossi	Deserranno
Setting	Mexican public-service recruitment	Village-promoter recruitment in Uganda
Variation	Recruitment financial incentives	Experimental earnings information
Outcomes	Applicant characteristics and recruitment outcomes	Beliefs and applicant selection
Lesson	Separate attraction from subsequent acceptance	Pay information can signal an unfamiliar job's nature

Different settings need not imply a universal rule that higher pay improves or worsens motivation. Applicant composition is an outcome, not a control to add mechanically.

Dal Bó, Finan, and Rossi (2013); Deserranno (2019).

More recruiting effort differs from more vacancies

- Davis–Faberman–Haltiwanger use establishment vacancy and hiring patterns to infer aspects of recruiting intensity.
- Carrillo-Tudela–Gartner–Kaas combine employer recruitment information, linked records, and a model of effort, standards, and wage policies.
- A residual, a reported effort measure, and an externally varied unit of capacity are different evidence objects.
- JOLTS openings are a stock at a point in time; hires are a flow over a month. Their ratio is not an application conversion rate.

Davis et al. (2013); Carrillo-Tudela et al. (2023); BLS JOLTS definitions. Inferred intensity is not directly recorded effort.

Direct effort measures sharpen the recruiting question

- **Setting:** Forsythe–Weinstein survey employers recruiting recent graduates.
- **Measurement:** planned recruiting effort and compensation, connected to hires and vacancies.
- **Evidence:** greater effort is linked to more hires per vacancy.
- **Limit:** reported effort is not random capacity assignment; difficult vacancies may induce additional effort.

Comparison to design

Distinguish planned effort, realized hours, applications reviewed, completed interviews, and eventual starts.

Forsythe and Weinstein (2025), *Recruiting Intensity, Hires and Vacancies: Evidence from Firm-Level Data*.

A capacity experiment must track displaced attention

- Prospectively randomize additional screening support across comparable requisitions or use a documented capacity shock with a defensible comparison.
- Follow queue length, reviewer hours, assessments, offers, withdrawals, and starts.
- Moving a recruiter toward treated vacancies can reduce attention to controls.
- Distinguish expanding capacity from reallocating a fixed resource.

Research proposal

Does extra capacity add information, prevent candidate loss while waiting, or both? These mechanisms need separate timing and outcome records.

Interview evaluations depend on the preceding candidate

- **Setting:** high-stakes admissions and hiring interviews.
- **Variation:** quasi-random candidate allocation and timing within actual pools, with independent quality assessments.
- **Outcome:** subjective evaluations respond to the preceding candidate's quality: a sequential contrast effect.
- **Limit:** scheduling assumptions and quality measurement matter. Proposed candidate reordering is evaluated through simulation, not a demonstrated field-policy success.

Radbruch and Schiprowski (2025), *Interview Sequences and the Formation of Subjective Assessments*. Published journal version.

Use an independent quality signal to study sequence

$$R_{ips} = \alpha_p + \lambda_s + \beta Q_i^{\text{ind}} + \gamma Q_{\text{prev}(i)}^{\text{ind}} + u_{ips}.$$

- R : focal evaluation; p : design-relevant pool; s : slot; α_p, λ_s : pool and slot effects.
- Q^{ind} : pre-existing or independently constructed quality signal; $\text{prev}(i)$: predecessor within the actual sequence.
- γ describes a predecessor-quality response. Causal interpretation needs justified order assignment and an uncontaminated signal.
- Alternating yes/no decisions can also reflect quotas, composition, or common shocks.

Teaching specification, not a transcription of every specification in Radbruch–Schiprowski (2025).

Reset the lag at the real session boundary

Session	Slot	Independent score	Correct lag	Wrong global lag
A	1	4	missing	missing
A	2	8	4	4
B	1	6	missing	8
B	2	5	6	6

- In a prospective extract, sort by evaluator, session, and timestamp; define tie and break rules. Existing studies may recover order without timestamps.
- The first interview has no predecessor in its session. Missing is not a zero-quality predecessor.
- Cancellations, substitutions, and missing assessments can alter the sequence relevant to the decision-maker.

Hypothetical four-interview illustration; these are not article data.

A lagged yes/no decision is not itself randomized

- ➊ Identify the scheduling pool and how order was assigned within it.
- ➋ Check rescheduling, no-shows, substitutions, and evaluator changes.
- ➌ Construct quality independently of the focal rating and respect information timing.
- ➍ Define the relevant predecessor and broader sequence composition.
- ➎ Match uncertainty to assignment and dependence, including evaluator/session relationships.

More candidate rows do not create more independent sequence assignments. Leave-out scores remove some mechanical dependence without randomizing evaluator assignment.

Discretion in Hiring compares tests with actual decisions

- **Setting:** job-test introduction and managerial decisions across service-sector firms.
- **Variation:** organizational rollout and departures from recommendations; not a randomized test-versus-manager experiment.
- **Outcome:** subsequent hire outcomes, centrally including tenure.
- **Finding:** average overrides in this setting do not appear to deliver superior private information reflected in those outcomes.
- **Limit:** tenure is not pure productivity; rejected candidates' performance remains unobserved.

Hoffman, Kahn, and Li (2018), *Discretion in Hiring*.

Observe the choice set before judging an override

Record	Why it matters
Available applicants	Who could actually receive the offer at that time?
Score and test version	What information and calibration did the manager see?
Displayed recommendation	Was the algorithm's advice available to the decision-maker?
Override and reason	What private information or inconsistency is proposed?
Offers, acceptance, starts	Did the preferred applicant join the organization?
Later outcomes and jobs	Did retention change through selection or job conditions?

A regression among hires cannot identify which rejected applicant should have received the offer.

Testing can affect several hiring outcomes at once

- **Setting:** Autor–Scarborough study retail establishments introducing job tests.
- **Variation:** test rollout, with its organizational comparison and assumptions.
- **Outcomes:** tenure and minority hiring, evaluated separately.
- **Lesson:** assess multiple consequences rather than assume a necessary quality–diversity tradeoff.
- **Limit:** no proprietary score becomes unbiased ground truth merely because it predicts one observed hire outcome.

Autor and Scarborough (2008), *Does Job Testing Harm Minority Workers? Evidence from Retail Establishments*.

A referral can change information, access, and support

- A referrer can identify a candidate, transmit private information, support the new hire, or participate in organizational decisions.
- Burks and coauthors compare referred and nonreferred workers across firms using productivity, turnover, and recruiting-cost outcomes.
- Turnover and recruiting costs can matter even when average measured productivity differences are limited.
- Realized referral status is not randomly assigned; these differences combine several mechanisms.

Burks, Cowgill, Hoffman, and Housman (2015), *The Value of Hiring through Employee Referrals*.

Referral experiments isolate particular channels

	Beaman–Magruder	Pallais–Sands
Setting	Social-network referral experiment	Field experiments on referrals
Variation	Referral incentives	Comparisons designed to separate channels
Objects	Who refers whom; information conveyed	Informational and other sources of referral value
Design lesson	The incentive changes selection into nomination	Map each contrast to the channel it varies

An experiment about one referral margin does not identify every mechanism behind a referred-worker retention or wage difference.

Beaman and Magruder (2012); Pallais and Sands (2016).

A referral program can affect people who were not referred

- **Setting:** an employee referral program introduced across stores.
- **Variation:** store-level randomized program assignment.
- **Outcomes:** direct referral outcomes and overall organizational effects, including retention among nonreferred employees.
- **Lesson:** a program's value is not exhausted by comparing referred hires with other hires.
- **Limit:** realized referral is a post-treatment status; inference must reflect assigned stores and spillovers.

Friebel, Heinz, Hoffman, and Zubanov (2023), *What Do Employee Referral Programs Do?*

Program assignment and referral status are different treatments

Program-level comparison

- All assigned stores
- Applicants and incumbent workers
- Direct effects and organizational responses
- Assignment-based interpretation

Referred-worker comparison

- Selected nominees or hires
- Referral status realized after assignment
- Misses nonreferred incumbents
- Needs its own identifying argument

Keep a separate referrer–candidate link table with nomination time, reward, eligibility, and coworker relationship.

Offers, acceptance, and starts run on different clocks

- Distinguish time since posting, application, interview, and offer. A one-day change in each clock alters different information.
- Horton randomizes employer recommendations in an online platform: an intermediary changes workers considered and hiring outcomes.
- A recommendation effect does not automatically establish full market welfare or the effect of a different offer deadline.
- Probation begins only for entrants. Define follow-up and censoring before comparing early exits.

Horton (2017); Dal Bó–Finan–Rossi (2013) separates recruiting offers from acceptance. Week 3 develops post-entry learning.

Hire-only means compare selected workforces

$$\begin{aligned} \mathbb{E}[Y_i^h(1) \mid H_i(1) = 1] - \mathbb{E}[Y_i^h(0) \mid H_i(0) = 1] \\ \neq \mathbb{E}[Y_i^h(1) - Y_i^h(0)] \quad \text{in general.} \end{aligned}$$

- z : policy regime; $H_i(z)$: potential hiring status; $Y_i^h(z)$: performance if employed in the specified job/regime.
- We observe performance only when the realized policy hires the worker.
- The selected-workforce mean is a useful policy outcome when labeled that way. It is not an individual productivity effect.
- Rejected candidates' latent job output is missing, not zero.

Two hiring rules change composition without changing anyone

Candidate	A	B	C	D	E	F
Fixed hypothetical output	10	9	8	7	6	5
Hired by Rule 0	yes	yes	no	no	no	no
Hired by Rule 1	yes	no	yes	no	no	no

$$\bar{Y}_0 = (10 + 9)/2 = 9.5, \quad \bar{Y}_1 = (10 + 8)/2 = 9.$$

- Both rules fill two slots. Selected total output changes from 19 to 18.
- No applicant's potential output changes; only the selected set changes.
- The instructor sees hypothetical truth; the analyst sees output only for those hired under the observed rule.

Constructed six-candidate example, not an empirical estimate or observed applicant dataset.


Constant hiring rates do not imply monotone selection

$$H_B(1) - H_B(0) = -1, \quad H_C(1) - H_C(0) = +1.$$

- One candidate leaves the hired set and another enters. Selection moves in both directions even though headcount is unchanged.
- Lee-style bounds require a defensible monotone-selection assumption in addition to the relevant assignment assumptions.
- Equal hiring rates alone do not justify applying a trimming procedure.
- If capacity also changes, compare total production and headcount costs as well as selected means.

Lee (2009), *Training, Wages, and Sample Selection*. The fixed-slot counterexample is hypothetical.

Evaluation choices and selection answer different questions



Constructed teaching examples — not empirical estimates. Panel B holds all six candidates' potential outputs fixed.

Information can change the decision; selection can change the mean.

Hypothetical teaching examples; reproducible figure from the displayed model and six-candidate table.

Hiring as Exploration makes learning an explicit objective

- **Setting:** resume screening and algorithmic exploration.
- **Evaluation:** retrospective policy comparison using historical decisions and labels.
- **Principal quality proxy:** eventual hiring among candidates considered for interviews, not universally observed job productivity.
- **Limit:** this is not a demonstrated prospective productivity RCT. Historical policies determine which labels exist.

Next design question

What would a prospective marginal invitation teach us, and which later outcomes would still remain missing?

Li, Raymond, and Bergman (2026), *Hiring as Exploration*, published journal version.

New observations help only for the outcomes they measure

- Pallais experimentally gives entry-level online workers a job and an evaluation, studying information and subsequent opportunities.
- A paid trial can reveal output on that task; it is not a complete measure of every future job or long-run match.
- Kleinberg and coauthors study prediction with selectively observed outcomes in bail decisions. Use it as a methodological comparison, not hiring evidence.
- Weighting or cross-fitting cannot create outcome support for groups never observed under relevant conditions.

Pallais (2014), *Inefficient Hiring in Entry-Level Labor Markets*; Kleinberg et al. (2018), *Human Decisions and Machine Predictions*.

A method name does not supply the recruiting comparison

Approach	Identifying information needed	Failure to investigate
Rollout / group-time DID	Credible untreated trends and treatment timing	Endogenous adoption; concurrent reforms
Test-cutoff RD	Threshold actually changes treatment	Manipulation or other discontinuities
Manager-leniency IV	Valid assignment and exclusion	Managers change training, tasks, or conditions too
Invitation weighting	Known probabilities and supported cells	Zero support or unobserved selection

Clustering handles specified dependence; it does not repair absent identifying variation.

Callaway and Sant'Anna (2021) addresses particular staggered-comparison problems, not endogenous adoption by itself.

Correspondence experiments identify contact-stage effects

- **Setting:** applications to large U.S. employers.
- **Variation:** randomized application characteristics.
- **Outcome:** employer contact, a defined entry-stage response.
- **Limit:** rejected applicants' job performance is not observed; correlations between organization and contact disparities are not causal organizational effects.

Scope of the claim

Explain the stage and intervention the experiment measures before extending its interpretation to wages, productivity, or firm organization.

Kline, Rose, and Walters (2022), *Systemic Discrimination Among Large U.S. Employers*.

Public documentation is not public applicant microdata

Resource	Usable object or route	Boundary
Applicant/test/referral logs	Event history under a specific agreement	No employer access assumed
IAB vacancy survey / linkage	Recruitment information and documentation	Microdata access is controlled
BLS JOLTS	Public aggregate openings and hires	No applicant screening sequence
Interview-sequence archive	Public-code pointer, construction information, moments	Original interviews proprietary
Pallais archive	Listed readmes and Stata programs	Listing does not prove all original data are public

Consult the chapter's access ledger before planning a project; no private-data application is required for the teaching lab.

Methods Box: six checks before estimating a recruiting effect

- ➊ **Risk set:** name the baseline cohort and eligible population.
- ➋ **Sequence:** retain event history and use the correct lag boundaries.
- ➌ **Information:** use what was known when the decision was made.
- ➍ **Comparison:** state why the intervention creates comparable cases.
- ➎ **Uncertainty:** count independent assignment units and shared slots.
- ➏ **Outcome availability:** show whose performance remains missing.

Request archived scores and recommendations, not overwritten final status. Keep unknown, pending, withdrawal, cancellation, and administrative censoring distinct.

Research Lab: reconstruct a design with synthetic sessions

- **Anchor:** Radbruch–Schiprowski's scheduling pools and independent-quality logic.
- **Offline mode:** synthetic design reconstruction; no proprietary interviews or published coefficients are reproduced.
- Generate randomized interview order, pre-interview scores, session effects, and a known sequential-contrast parameter.
- Construct within-session predecessor scores, estimate with session fixed effects, and compute session-clustered uncertainty. This teaching DGP sets the earlier specification's slot effects $\lambda_s = 0$.

Interpret the success correctly

Recovering the known simulated parameter checks the exercise. It does not establish the article's empirical scheduling assumptions.

Diagnose boundaries, labels, and fixed-slot selection

- Compare the valid lag with an incorrect global lag that crosses sessions.
- Reject non-nested funnel counts; separately diagnose whether stages use compatible cohorts.
- Apply both hiring rules to the same six candidates and count entries into and exits from the selected set.
- Keep rejected candidates' synthetic potential output out of the analyst file; instructor truth is a separate synthetic exercise key.
- Explain why constant hiring rates fail to establish monotone selection.

All generated observations and estimates are synthetic. The smoke executes code checks offline; cohort comparability still needs a design argument.

Frontier proposals add a specific missing observation

Proposed direction	Observation and comparison to add
Capacity versus candidate loss	Vary reviewer support; observe queues, withdrawal, offers, starts, and displaced attention.
Exploration with outcomes	Randomize marginal paid invitations; retain selection probabilities, standardized output, and worker follow-up.
Stage-specific entry barriers	Connect posted rules, exposure, applications, and employer responses; distinguish who applies from who is contacted.

Feng et al. (September 2026 working paper) motivates the third direction with age restrictions and platform stages. Inspect the actual design before calling it an RD.

Research proposals, not claims of unstudied gaps. Feng et al., *Signals, Steps, and Setbacks*, NBER Working Paper 35709.

Three required readings, with rotating frontier presentations

Reading	Extract before discussing the finding
Hoffman–Kahn–Li (2018)	Available choice set, recommendation, manager decision, later outcome
Radbruch–Schiprowski (2025)	Scheduling block, independent quality, sequence comparison
Friebel et al. (2023)	Program assignment, all-store outcomes, realized referral status
Rotating: Li–Raymond–Bergman (2026)	Historical labels, retrospective evaluation, prospective learning
Rotating: Forsythe–Weinstein (2025)	Reported effort versus inferred intensity

Use the chapter's pathways for theory, channels, referral mechanisms, and selection methods. The full bibliography is a research menu.

Transfer: specify the recruiting study you would run

Two-page design memo

Rule change; baseline risk set; stage ledger; counterfactual; missing performance; assignment and inference unit; data access; economically meaningful limitation.

- Use *Discretion in Hiring* or *Hiring as Exploration* to propose one bounded comparison.
- Separate an experimental or quasi-experimental estimate from model-based extrapolation.
- Explain how the design changes if the firm has two fixed slots.
- State a worker outcome beyond the firm's tenure or fill-rate objective.

Entry shapes the information the firm will later observe

- The recruiting funnel links costly employer choices to applicant responses.
- Information, attention, composition, and acceptance can change together; the design must say which margin it compares.
- Observing hires alone does not reveal every rejected candidate's performance.

Bridge to Week 3

Once workers enter, what do output, quality, ratings, tasks, and supervisors teach the employer? Carry forward the distinction between changing performance and changing who is observed.