Early-Career Discrimination:

Spiraling or Self-Correcting?

Arjada Bardhi[1] Yingni Guo[2] Bruno Strulovici[3]

SaMMF Workshop: Theories of Discrimination
June 2020

[1]Duke [23]Northwestern

Early-career discrimination

- Employers know little about the true productivity of early-career workers
- To address this, they rely on proxies for workers' productivities
  - observable characteristics (race, gender, ethnicity etc.)
  - Goldin and Rouse (2000), Pager (2003), Bertrand and Mullainathan (2004), Bertrand and Duflo (2016) etc.
  - once hired, on-the-job performance is informative
- When jobs are scarce, groups that are discriminated at the start miss on early opportunities
- Even if groups' productivity distributions are very similar!

Questions

1. How important is early-career discrimination for workers' lifetime prospects?
2. As groups' productivities converge, do their payoffs converge too?

Two conjectures on the impact of group belonging:

1. small difference $\rightarrow$ employers learn $\rightarrow$ errors in hiring corrected quickly $\rightarrow$ little impact
2. small difference $\rightarrow$ unequal early career opportunities $\rightarrow$ different career trajectories $\rightarrow$ significant impact

Key insight:

How employers learn about workers' productivity matters.

Baseline model

Model

Players and types

- One employer and two workers: a and b
- Each worker from a distinct social group
- Productivity type of worker i is either high or low: $\theta_i \in \{h, \ell\}$
- Group i's average productivity: $p_i := \Pr(\theta_i = h)$

Comparable social groups

(i) group a has higher productivity: $p_a > p_b$

(ii) groups have almost identical productivity distributions: $p_b \to p_a$

Model

Task allocation and payoffs

- Continuous time $t \in [0, \infty)$ and long-lived players
- At each t, employer allocates a divisible task

$$\{\text{ worker } a, \text{ worker } b, \text{ safe arm }\}$$

- Employer's flow payoff:
  - $v > 0$ if task goes to a high-type worker
  - 0 if task goes to a low-type worker
  - $s \in (0, v)$ if safe arm
- Worker's flow payoff:
  - fixed wage $w = 1$ if allocated the task
  - 0 otherwise

Employer's problem is a standard three-armed bandit problem.

Model

Learning environments

worker i is allocated the task over $[t, t + \mathrm{d}t)$

$$\Downarrow$$

employer learns about θ_i over $[t, t + \mathrm{d}t)$

We contrast two learning environments:

Breakthrough

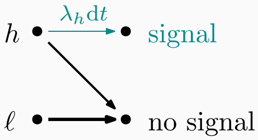

Breakdown

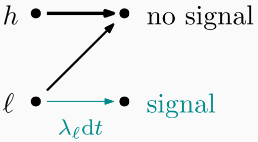

Poisson signals arrive at rate λ_h and λ_ℓ respectively.

Interpreting learning environments

- Learning varies by occupation, rank, etc

Interpreting learning environments

- Learning varies by occupation, rank, etc
- Tracking under-performance (breakdowns) vs. over-performance (breakthroughs)

Interpreting learning environments

- Learning varies by occupation, rank, etc
- Tracking under-performance (breakdowns) vs. over-performance (breakthroughs)
- Jacobs (1981), Baron and Kreps (1999):

“star jobs” vs. “guardian jobs”

Interpreting learning environments

- Learning varies by occupation, rank, etc
- Tracking under-performance (breakdowns) vs. over-performance (breakthroughs)
- Jacobs (1981), Baron and Kreps (1999):

"star jobs" vs. "guardian jobs"

> ' *The first-rate salesman can often add a significant increment to the performance of his organization while his inferior will not impose unacceptable costs.'* Jacobs, 1981

Interpreting learning environments

- Learning varies by occupation, rank, etc
- Tracking under-performance (breakdowns) vs. over-performance (breakthroughs)
- Jacobs (1981), Baron and Kreps (1999):

"star jobs" vs. "guardian jobs"

> '*The airline pilot who misses a landing or the operative who inadvertently blocks a long assembly line will produce rather destructive effects, but an outstanding performance in either position will be of little consequence for the organization.*'
>
> Jacobs, 1981

Baron and Kreps (1999)

Breakthrough vs. breakdown learning

1. Do workers' lifetime payoffs converge as $p_b \uparrow p_a$?
2. Which learning environment, if any, grants a disproportionate first-hire advantage?

Related work

1. **Statistical discrimination**:
   - Phelps (1972), Aigner and Cain (1977), Cornell and Welch (1996), Fershtman and Pavan (2020)
   - Arrow (1973), Foster and Vohra (1992), Coate and Loury (1993), Moro and Norman (2004)

   Cumulative discrimination: Blank, Dabady, and Citro (2004), Blank (2005)

   Discrimination in hiring and referrals: Bertrand and Mullainathan (2004), Bertrand and Duflo (2017), Sarsons (2019)

2. **Employer learning**: Farber and Gibbons (1996), Altonji and Pierret (2001), Altonji (2005), Lange (2007), Antonovics and Golan (2012), Mansour (2012), Bose and Lang (2017)

3. **Bandit approach**: Felli and Harris (1996), Bergemann and Valimaki (1996), Keller, Rady, and Cripps (2005), Strulovici (2010), Keller and Rady (2010, 2015)

A stark contrast

Breakthrough learning

Optimal allocation

At each t, employer allocates the task to either the worker that is more likely to be h or the safe arm ($\underline{p}$).

Breakthrough learning

Optimal allocation

At each t, employer allocates the task to either the worker that is more likely to be h or the safe arm ($\underline{p}$).

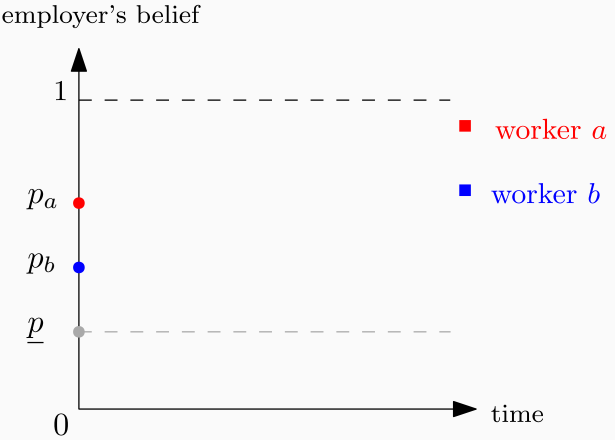

Breakthrough learning

Optimal allocation

At each t, employer allocates the task to either the worker that is more likely to be h or the safe arm ($\underline{p}$).

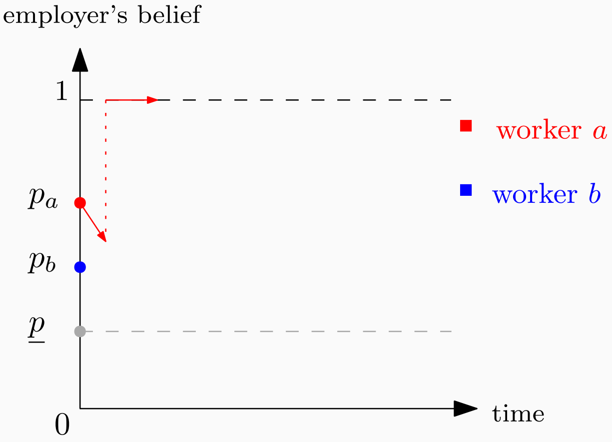

Breakthrough learning

Optimal allocation

At each t, employer allocates the task to either the worker that is more likely to be h or the safe arm ($\underline{p}$).

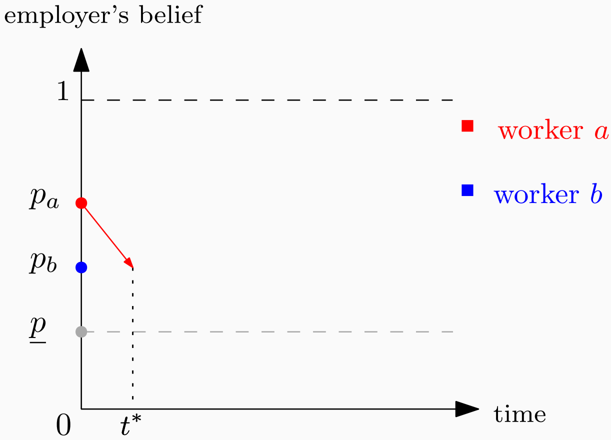

Breakthrough learning

Optimal allocation

At each t, employer allocates the task to either the worker that is more likely to be h or the safe arm ($\underline{p}$).

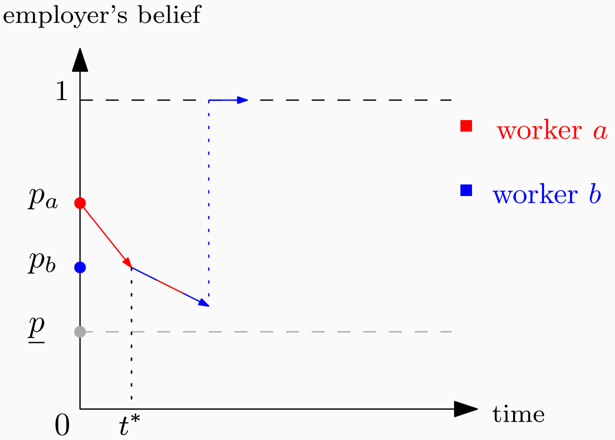

Breakthrough learning

Optimal allocation

At each t, employer allocates the task to either the worker that is more likely to be h or the safe arm ($\underline{p}$).

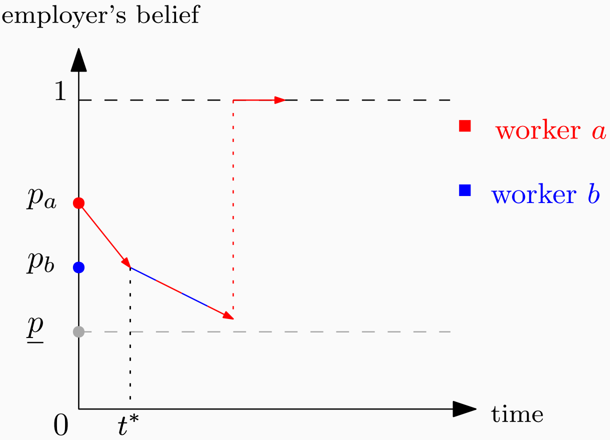

Self-correction under breakthroughs

Proposition 1a

As $p_b \uparrow p_a$, the expected payoff of worker b converges to that of worker a.

- task assigned exclusively to worker a over $[0, t^*]$

$$t^* = \frac{1}{\lambda_h} \log \left(\frac{p_a/(1-p_a)}{p_b/(1-p_b)} \right)$$

- workers treated symmetrically after t^*
- as $p_b \uparrow p_a$, grace period $t^* \to 0$
- the advantage of worker a vanishes

Breakdown learning

Optimal allocation

At each t, employer allocates the task to either the worker that is more likely to be h or the safe arm ($\underline{p}$).

Breakdown learning

Optimal allocation

At each t, employer allocates the task to either the worker that is more likely to be h or the safe arm ($\underline{p}$).

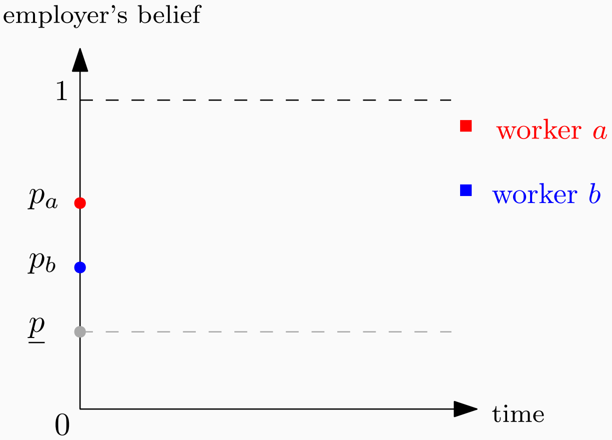

Breakdown learning

Optimal allocation

At each t, employer allocates the task to either the worker that is more likely to be h or the safe arm ($\underline{p}$).

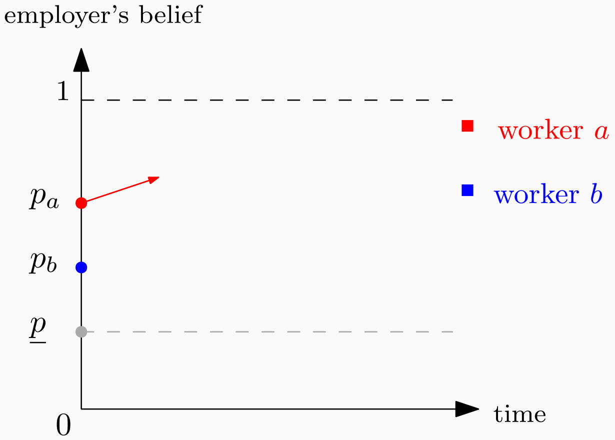

Breakdown learning

Optimal allocation

At each t, employer allocates the task to either the worker that is more likely to be h or the safe arm ($\underline{p}$).

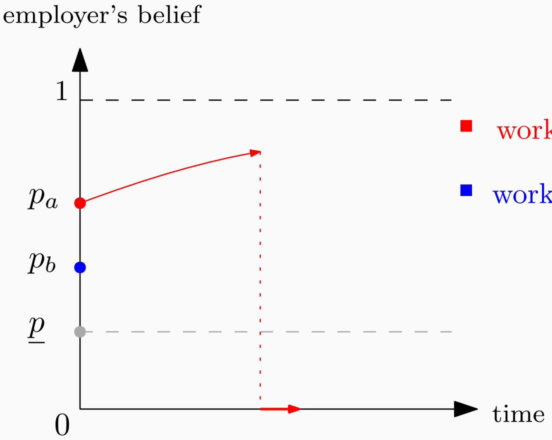

Breakdown learning

Optimal allocation

At each t, employer allocates the task to either the worker that is more likely to be h or the safe arm ($\underline{p}$).

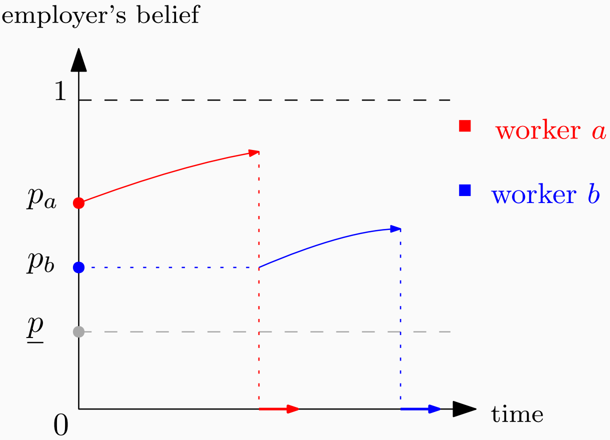

Spiraling under breakdowns

Proposition 1b

As $p_b \uparrow p_a$, the ratio of the expected payoff of worker b to that of worker a converges to

$$(1 - p_a)\frac{\lambda_\ell}{\lambda_\ell + r} < 1.$$

- task assigned to worker a until he realizes a breakdown
- worker a's payoff

$$\underbrace{p_a}_{\text{no breakdown ever}} + (1 - p_a) \cdot \underbrace{\frac{r}{\lambda_\ell + r}}_{\text{expected time until a breakdown}}$$

Spiraling under breakdowns

Proposition 1b

As $p_b \uparrow p_a$, the ratio of the expected payoff of worker b to that of worker a converges to

$$(1 - p_a)\frac{\lambda_\ell}{\lambda_\ell + r} < 1.$$

- task assigned to worker a until he realizes a breakdown
- worker a's payoff

$$\underbrace{p_a}_{\text{no breakdown ever}} + (1 - p_a) \cdot \underbrace{\frac{r}{\lambda_\ell + r}}_{\text{expected time until a breakdown}}$$

- worker b's payoff

$$\underbrace{(1 - p_a)\frac{\lambda_\ell}{\lambda_\ell + r}}_{b \text{ gets a chance}} \left(p_b + (1 - p_b)\frac{r}{\lambda_\ell + r}\right)$$

Contrast between breakthrough and breakdown learning

As $p_b \uparrow p_a$, worker a's advantage from early-career discrimination:

- vanishes under breakthrough learning
  - comparable workers $\Rightarrow$ comparable lifetime payoffs
- persists under breakdown learning
  - comparable workers $\not\Rightarrow$ comparable lifetime payoffs
  - even for very fast learning: $\lambda_\ell \to +\infty$

Extensions

We explore this contrast in three directions:

(i) Large labor market

(ii) Flexible wages

(iii) Opportunity to invest in productivity

- Inequality even higher in the breakdown environment!

The contrast is moreover robust to:

(iv) Misspecified beliefs by employer: $p_a = p_b$ but $\tilde{p}_a > \tilde{p}_b$

(v) Inconclusive breakthroughs / breakdowns

(vi) Group differences in speed of learning: $\lambda^b \uparrow \lambda^a$

Investment in productivity

Large labor market

Large market

Breakthroughs

- unit mass of tasks, α mass *a*-workers, β mass *b*-workers
- frictionless matching
- task scarcity: more workers than tasks

Large market

Breakthroughs

- unit mass of tasks, α mass a-workers, β mass b-workers
- frictionless matching
- task scarcity: more workers than tasks

Under breakthrough learning:

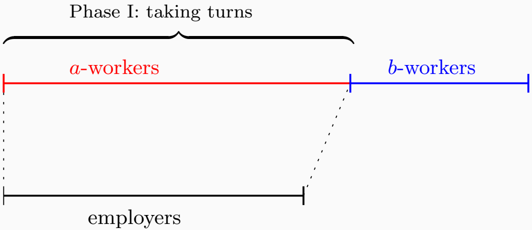

Large market

Breakthroughs

- unit mass of tasks, α mass a-workers, β mass b-workers
- frictionless matching
- task scarcity: more workers than tasks

Under breakthrough learning:

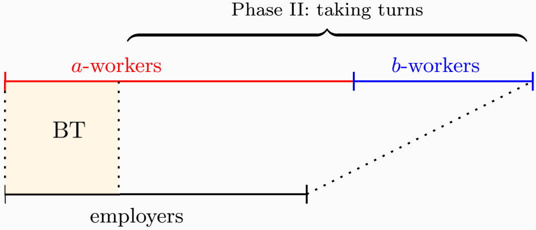

Large market

Breakthroughs

- unit mass of tasks, α mass *a*-workers, β mass *b*-workers
- frictionless matching
- task scarcity: more workers than tasks

Under breakthrough learning:

Phase I : tasks split between *a*-workers only

Phase II : remaining tasks split between *a*-workers and all *b*-workers

Large market

Breakthroughs

- unit mass of tasks, α mass *a*-workers, β mass *b*-workers
- frictionless matching
- task scarcity: more workers than tasks

Under breakthrough learning:

Phase I : tasks split between *a*-workers only

Phase II : remaining tasks split between *a*-workers and all *b*-workers

Self-correction under breakthroughs

Delay for group *b* vanishes as $p_b \uparrow p_a$.

Large market

Breakdowns

- unit mass of tasks, α mass a-workers, β mass b-workers
- frictionless matching
- task scarcity: more workers than tasks

Large market

Breakdowns

- unit mass of tasks, α mass a-workers, β mass b-workers
- frictionless matching
- task scarcity: more workers than tasks

Under breakdown learning:

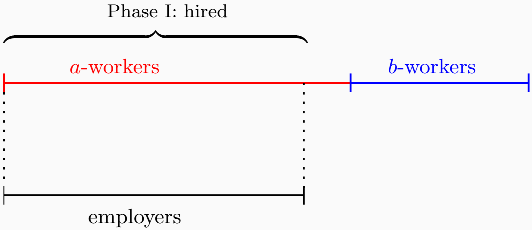

Large market

Breakdowns

- unit mass of tasks, α mass a-workers, β mass b-workers
- frictionless matching
- task scarcity: more workers than tasks

Under breakdown learning:

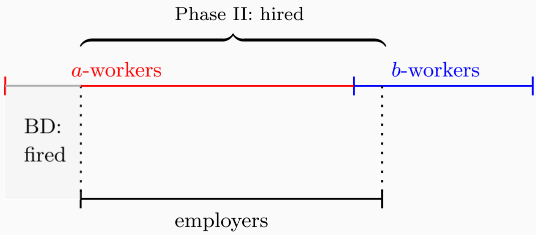

Large market

Breakdowns

- unit mass of tasks, α mass *a*-workers, β mass *b*-workers
- frictionless matching
- task scarcity: more workers than tasks

Under breakdown learning:

Phase I : *a*-workers hired only

Phase II : *b*-workers hired after sufficiently many *a*-workers failed

Large market

Breakdowns

- unit mass of tasks, α mass a-workers, β mass b-workers
- frictionless matching
- task scarcity: more workers than tasks

Under breakdown learning:

Phase I : a-workers hired only

Phase II : b-workers hired after sufficiently many a-workers failed

Spiraling under breakdowns

Delay for group b does not vanish as $p_b \uparrow p_a$.

How does group size affect inequality?

Proposition (Inequality increases in task scarcity)

Let $\alpha > 1$ and $\beta > 0$. As $p_b \uparrow p_a$, the limiting ratio of the expected payoff of a b-worker to that of an a-worker decreases *in both α and β.*

tasks become scarcer

$\Rightarrow$ more competition among workers

$\Rightarrow$ *b*-workers are hurt more than *a*-workers

$\Rightarrow$ inequality deepens

While all groups suffer during economic downturns, some suffer disproportionately more.

Flexible wages

Can flexible wages fix spiraling?

Answer: No, as long as wages are non-negative.

Approach:

- cooperative solution: dynamic stability as in Ali and Liu (2020)
- repeating any stable stage-game matching (Shapley and Shubik, 1971) is dynamically stable

Can flexible wages fix spiraling?

Solution:

- workers with the highest belief are matched at any instant
- there is a history-dependent marginal belief p^M
- wage schedule is convex
  - $(p - p^M)\,v$ for matched workers and 0 for unmatched ones

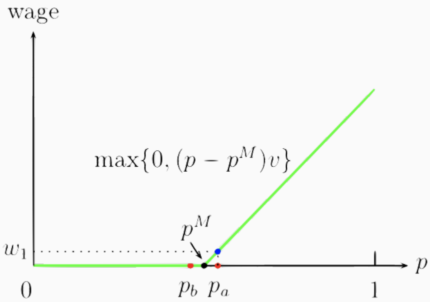

Can flexible wages fix spiraling?

Intuition

- more learning about a worker's type $\Rightarrow$ higher expected wage
- delay for group b does not vanish as $p_b \uparrow p_a$
- more is learned about a-workers than b-workers

Two-period intuition:

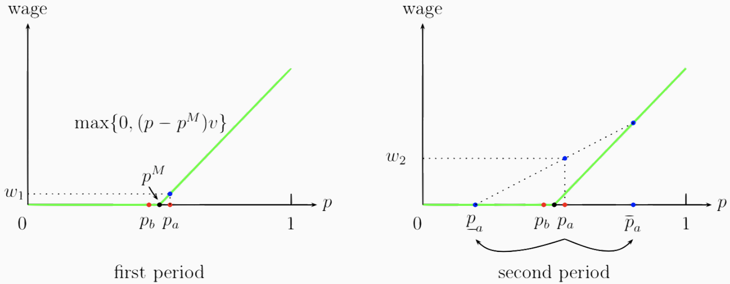

Final thoughts

Final thoughts

> ‘*How economically relevant statistical discrimination is depends on how fast employers learn about workers’ productive types.*’ Lange (2007)

- The nature of learning – not just the speed – is key for early-career discrimination.
- Early-career discrimination among comparable workers can have a significant lifetime impact
- More empirical work needed on the persistence and magnitude of discrimination in star vs. guardian jobs

Thank you.

Interpreting learning environments

Adapted from Fig. 2-2 in Baron and Kreps (1999)

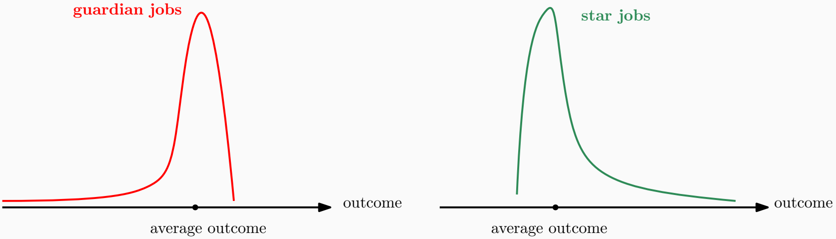

Figure 1: Distribution of outcomes for different types of jobs.

Back

Investment in productivity

How we model the investment opportunity?

- Before $t = 0$, each ℓ-type worker draws his investment cost from distribution F on $[0, 1]$, and decides whether to invest
- If a low-type worker invests, his type improves to h
- The pre-investment and post-investment types are private information to the worker
- F is the same for both groups

Investment in productivity

What is common between environments?

- (Post-investment) favored worker has stronger incentives to invest than the discriminated one
- This self-fulfilling force leads to multiple equilibria
- There exist equilibria in which *b* overtakes *a* and becomes favored

Equilibrium sets

We compare the equilibrium sets across two learning environments.

Investment in productivity

Preview of key results

Result 1: equilibrium payoff

- Investment does not disturb the self-correcting property of breakthroughs
- Investment exacerbates spiraling under breakdowns: it makes the workers' payoffs more unequal than without investment

Result 2: investment behavior

- When learning is sufficiently fast, breakdown learning leads to more polarized investment across the two workers than breakthrough learning does

Back