Early-Career Discrimination: Spiraling or Self-Correcting?

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

[1]Duke [23]Northwestern

Virtual MD Seminar, Sep 2020

Motivating setting

Medical referrals (Sarsons, 2019)

- Male and female surgeons compete for referrals from physicians
- Physicians gather new information about a surgeon's ability only if the surgeon performs a surgery
- Male and female surgeons have comparable abilities

> *'Women have a lower average ability and a slightly lower variance of ability, but the differences are small.'* Sarsons (2019)

FIGURE 2: DISTRIBUTION OF SURGEON ABILITY

(a) Matched Sample

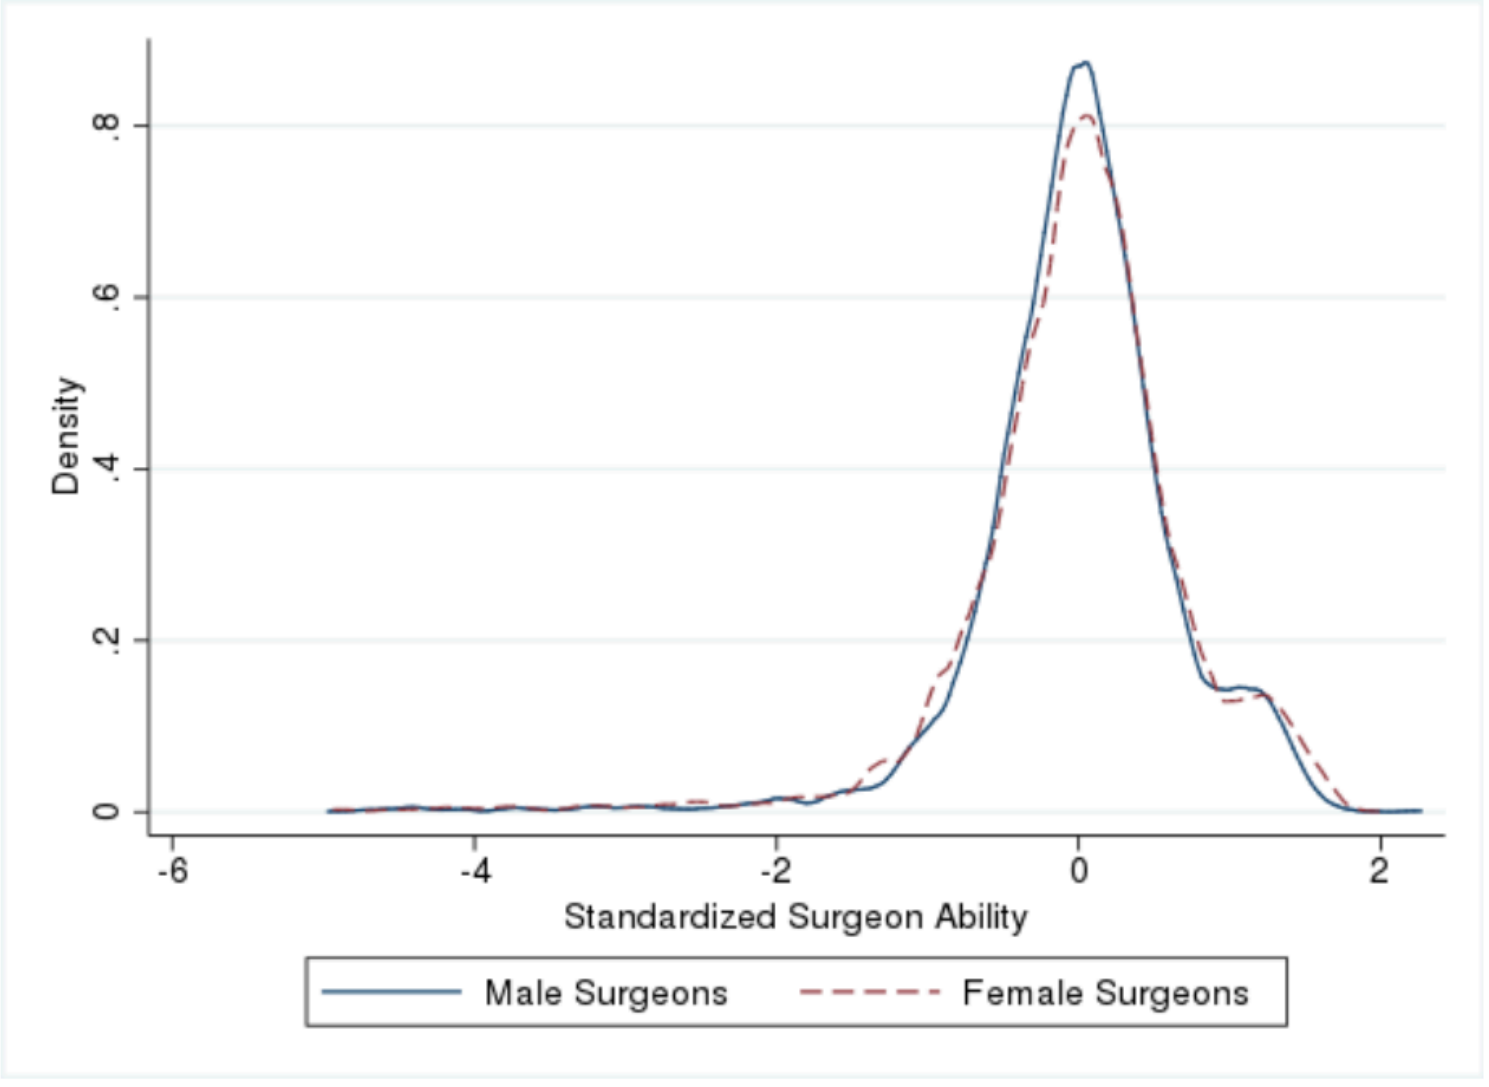

Motivating setting

Generally,

- workers from different social groups compete for tasks
- employers learn about a worker's productivity only if the worker performs a task

  today's belief $\Rightarrow$ today's task allocation
  $\Rightarrow$ tomorrow's belief $\Rightarrow$ tomorrow's task allocation
  $\Rightarrow$ the day after tomorrow's belief $\Rightarrow$ ...
- groups have comparable productivity distributions

Questions

- How does workers' group belonging (gender, race, etc) affect their lifetime payoffs?
- When workers are young, employers use group belonging to infer how productive they are

  But what happens in the long run?

  Does the impact of such early-career discrimination vanish or intensify over time?
- As groups' productivity distributions converge, do their payoffs converge too?

Two opposite conjectures

Group belonging has little impact

- groups have comparable productivity distributions
- employers get chances to learn about workers' productivity

Group belonging has significant impact

- opportunities to perform tasks matter
- without those early opportunities, it's hard to move up the career ladder

Preview of key results (I)

- The answer depends on how employers learn
- Certain learning environments deliver comparable payoffs to comparable groups (self-correcting)
- Other learning environments translate small prior differences into large payoff disparities across groups (spiraling)

- Self-correcting environments are those that track successes
- Spiraling environments are those that track failures

Preview of key results (II)

- This contrast persists with both fixed and flexible wages
- In a spiraling environment, comparable groups face very different wage paths

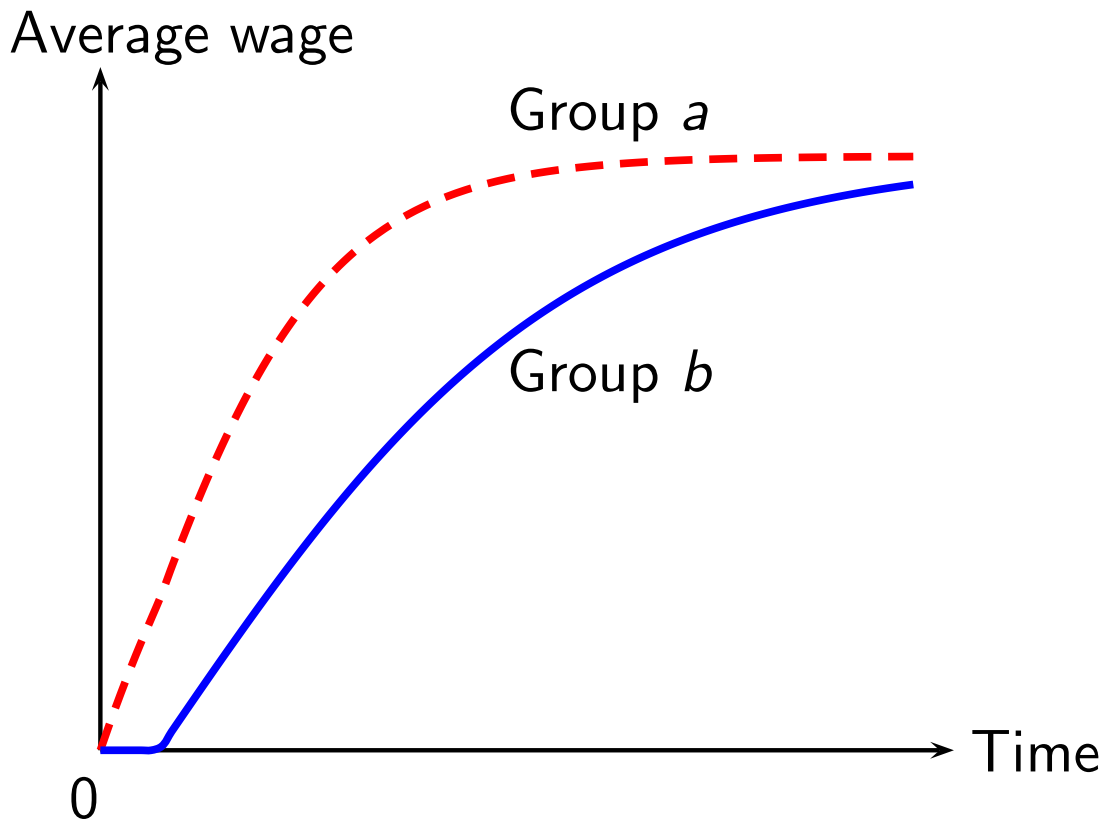

Preview of key results (II)

- Statement on Gender Salary Equity by *Association of Women Surgeons* in 2017:

  'The disparities women face in compensation at entry level positions lead to a persistent trend of unequal pay for equal work throughout the course of their careers.'

- Arcidiacono, Bayer and Hizmo (2010) document that racial wage gaps are small at early career stages but widen with labor market experience

Preview of key results (III)

The contrast persists when workers can invest in their productivity

- Spiraling environments polarize incentives to invest across groups
- Self-correcting environments lead to more equalized incentives to invest across groups
- If learning is sufficiently fast, employers prefer spiraling environments

  Tradeoff: *efficiency* for employers versus *equality* between the workers

Related work

- **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**: Goldin and Rouse (2000), Bertrand and Mullainathan (2004), Bertrand and Duflo (2017), Sarsons (2019)

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

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

Roadmap

- Baseline model
- Self-correcting vs spiraling
- Large labor markets
- Flexible wages
- Investment in productivity

Players and types

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

- Worker a is ex-ante more productive, but workers are comparable:

$$p_b < p_a, \text{ but } p_b \uparrow p_a$$

Task allocation

- Every day, the employer has a task to allocate
- He gets $v > 0$, if task goes to a worker of high type
- He gets 0, if task goes to a worker of low type

- A worker gets $w = 1$ if he gets the task that day (fixed wage)
- He gets 0 otherwise

Learning by allocating

- Learn about worker i's productivity only if i performs a task
- Breakthrough learning:
  - If task is performed by a low-type worker, no signal
  - If performed by a high-type worker, a breakthrough occurs sometimes
  - Academia jobs/R&D
- Breakdown learning:
  - If task is performed by a high-type worker, no signal
  - If performed by a low-type worker, a breakdown occurs sometimes

Interpreting breakthrough vs breakdown learning

- Intrinsic feature of the job considered
- 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.[...] The novice salesman is given only a limited time to produce. The result is that there tends to be a continuously rotating pool of newcomers who stay with the organization for short periods of time, while those who manage to be successful receive large rewards and some guarantee of future security'

Jacobs (1981)

Interpreting breakthrough vs breakdown learning

- Intrinsic feature of the job considered
- 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)

Model summary

- Continuous time $t \in [0, \infty)$, discount rate $r > 0$
- The employer faces a standard bandit problem
  - Workers a, b are two bandit arms with priors $p_b < p_a \in (0, 1)$
  - At each t, allocate the task to whoever is more likely to have high type
  - If both look too unproductive (below p_s), assign task to a safe arm
- Breakthrough: if task goes to h, breakthrough occurs at Poisson rate λ_h
- Breakdown: if task goes to ℓ, breakdown occurs at Poisson rate λ_ℓ

Roadmap

- Baseline model
- Self-correcting vs spiraling
- Large labor markets
- Flexible wages
- Investment in productivity

Breakthrough learning

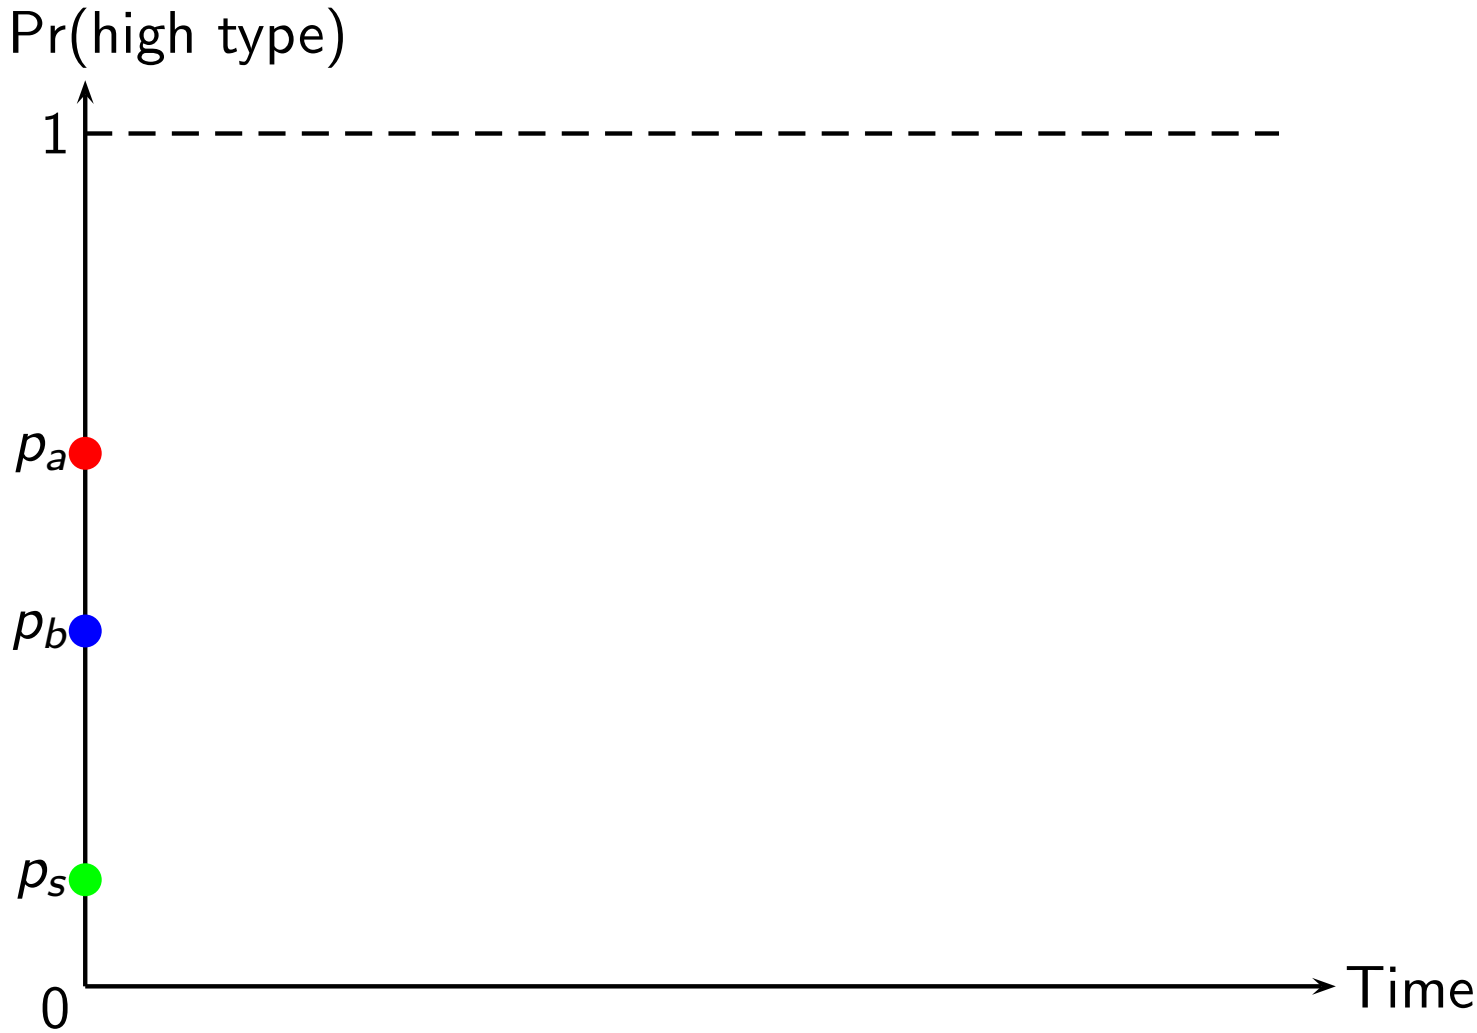

Breakthrough learning

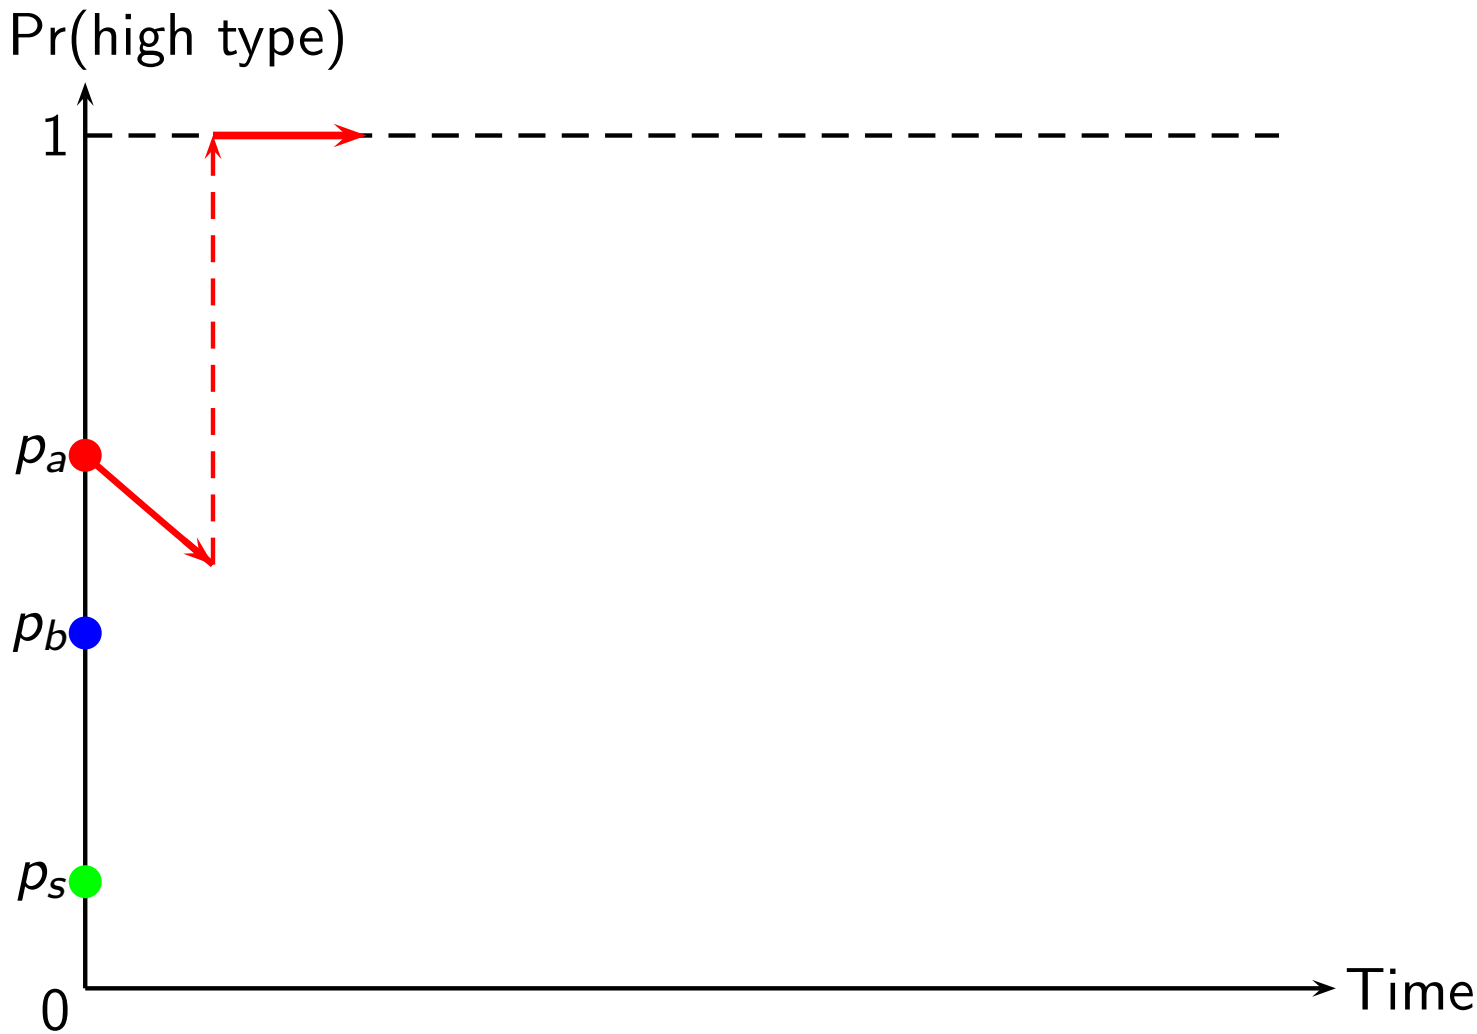

Breakthrough learning

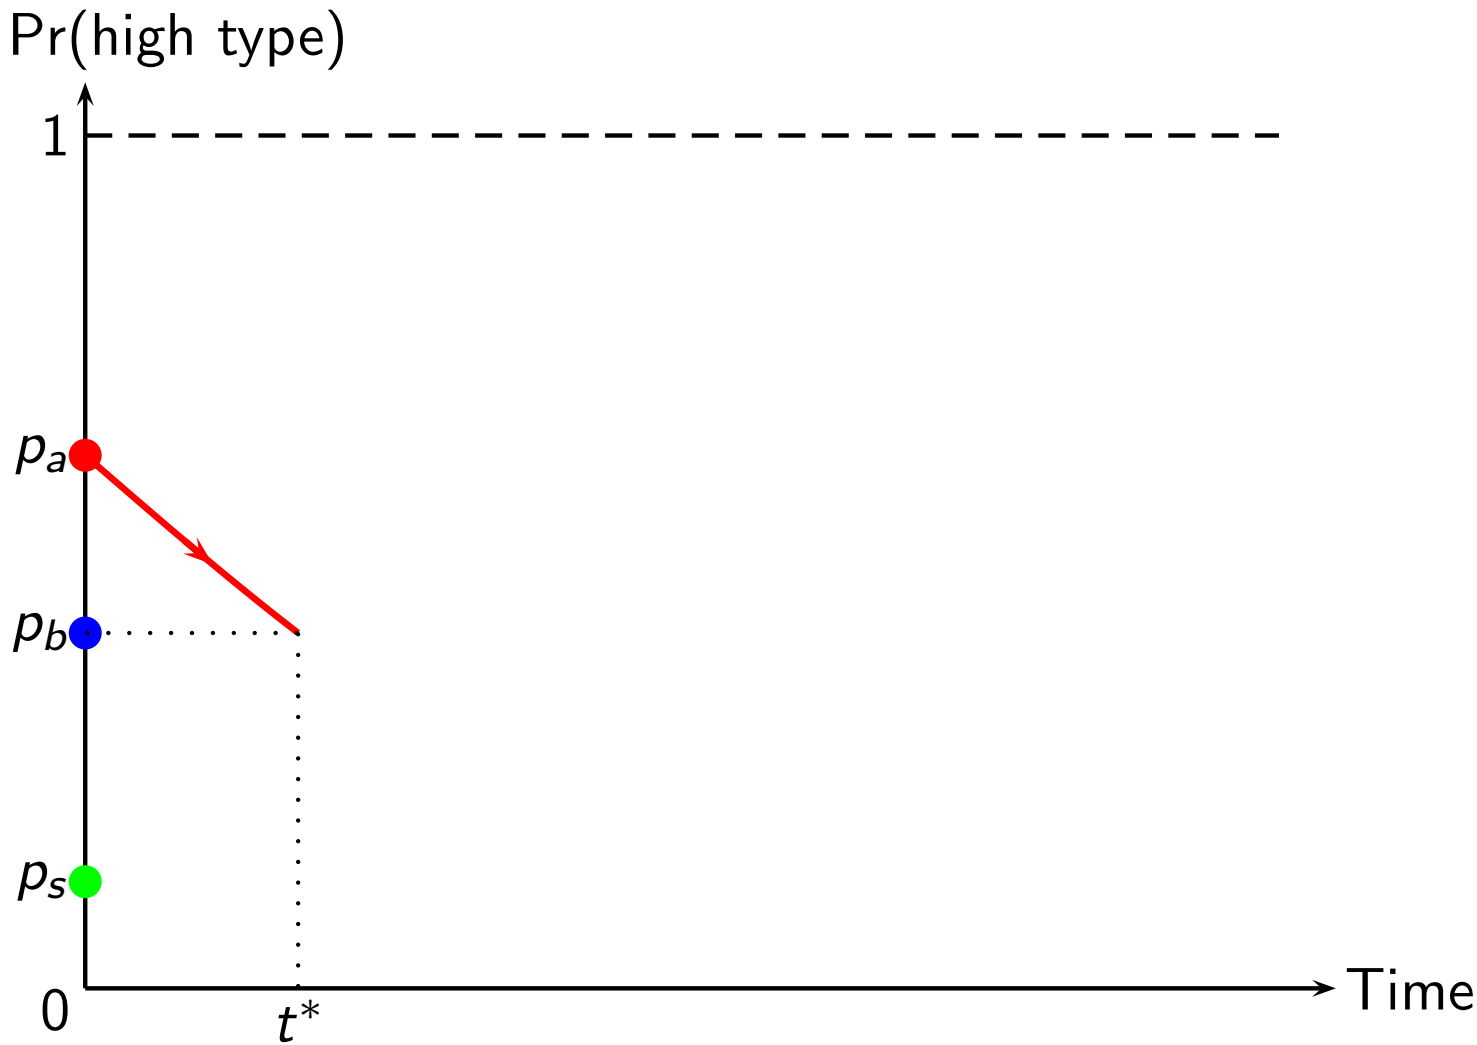

Breakthrough learning

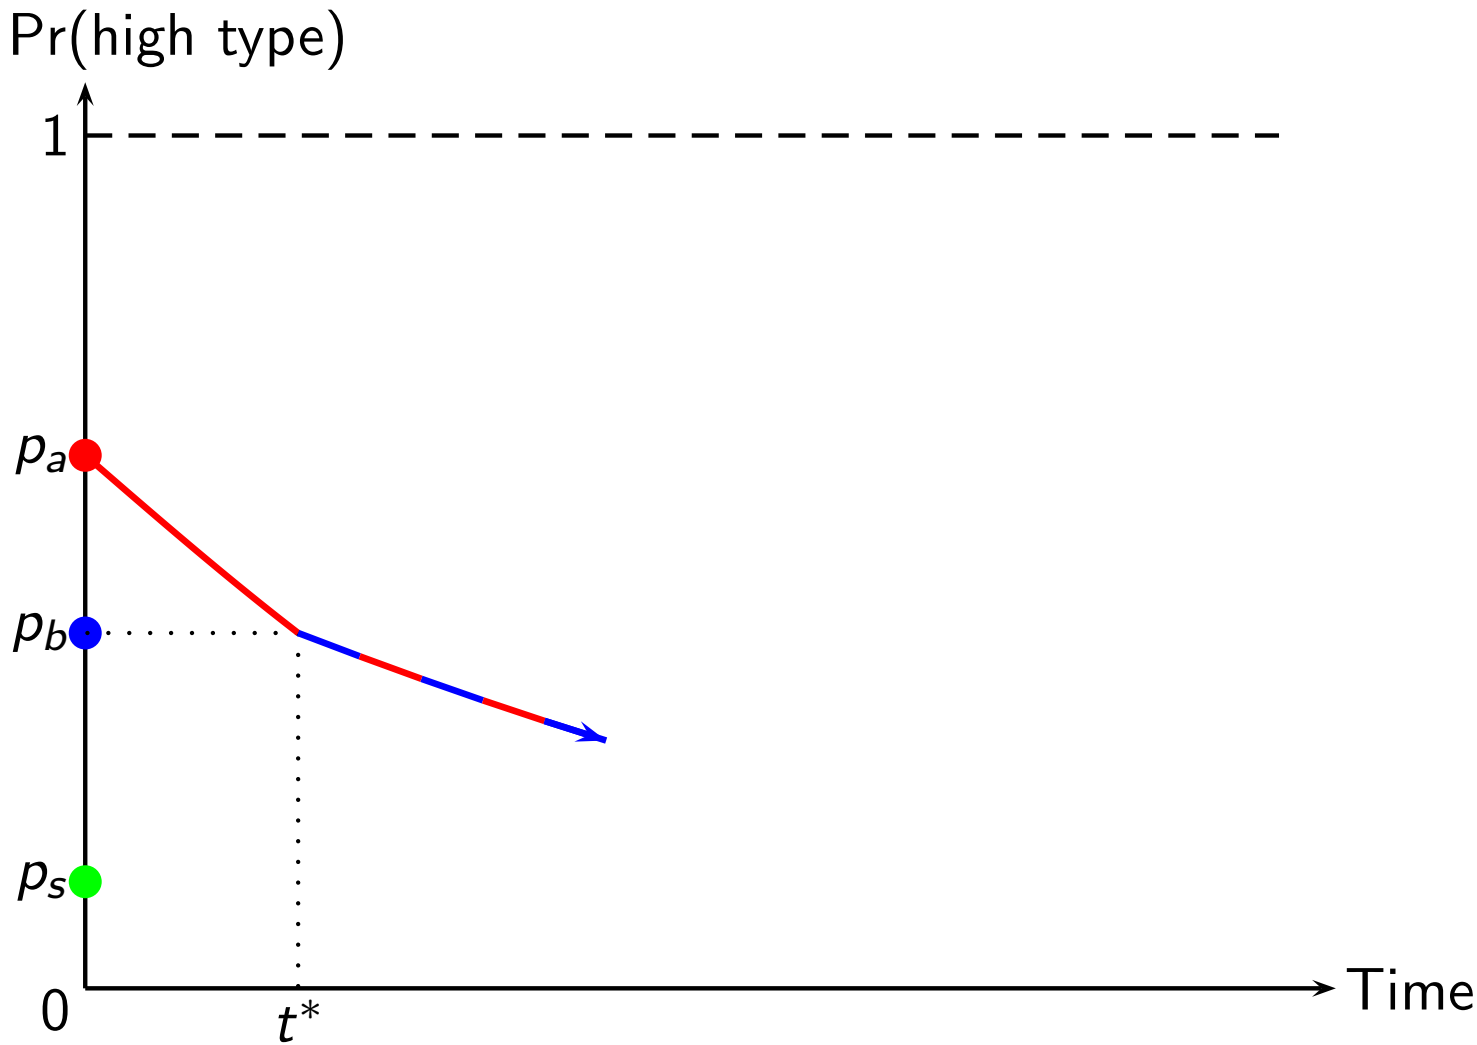

Breakthrough learning

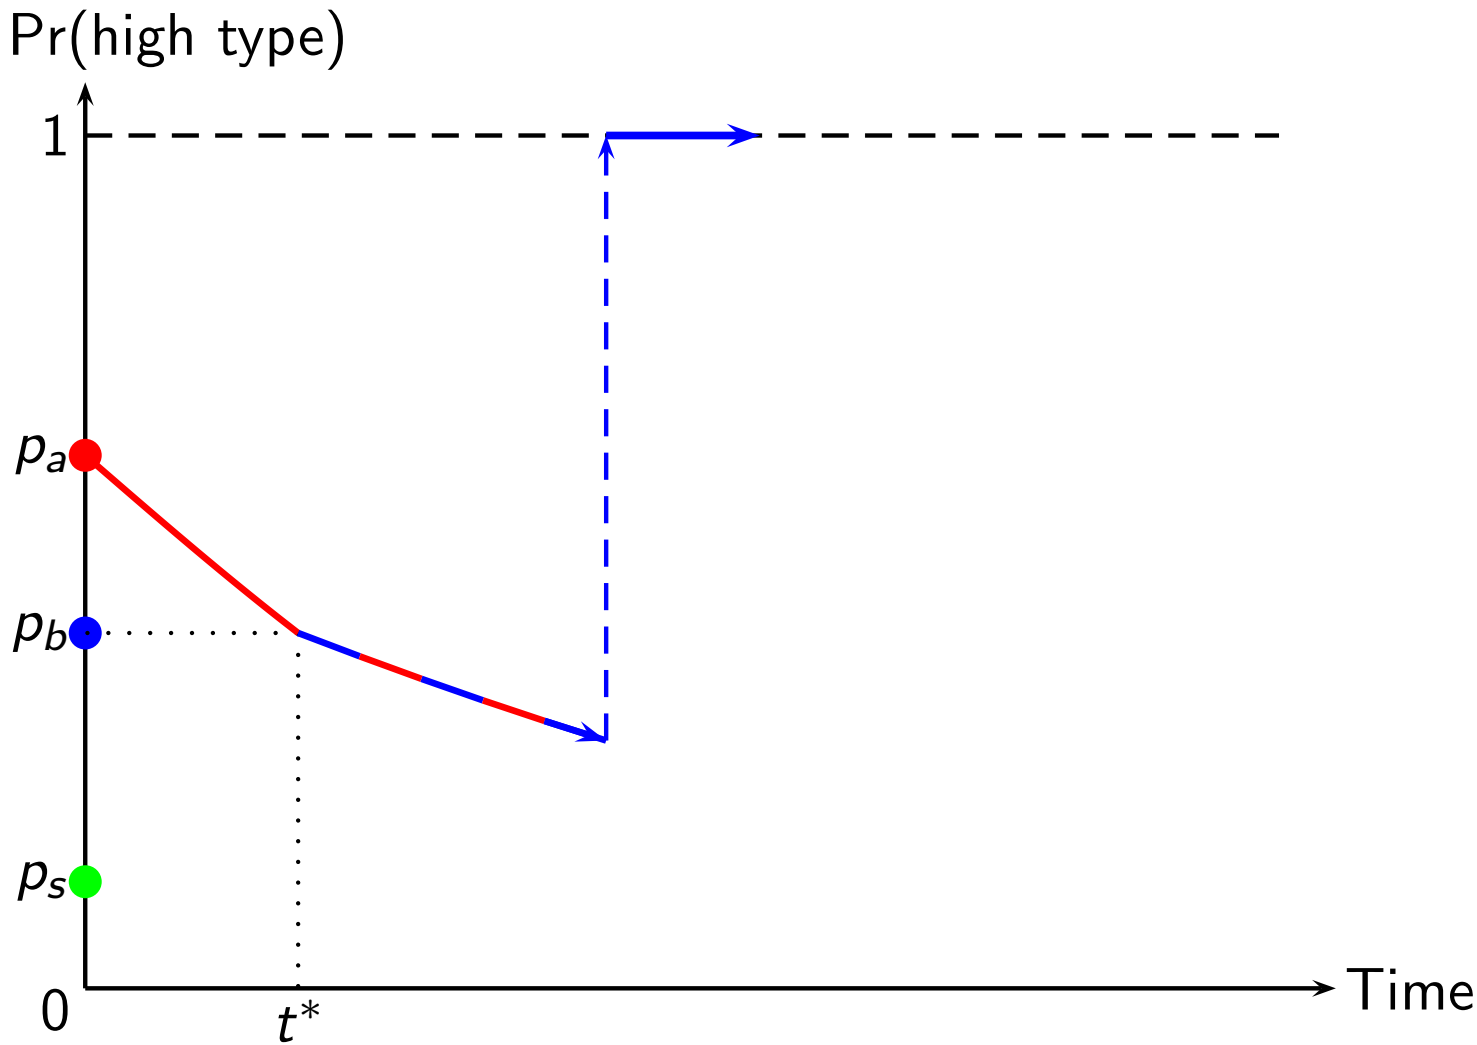

Breakthrough learning

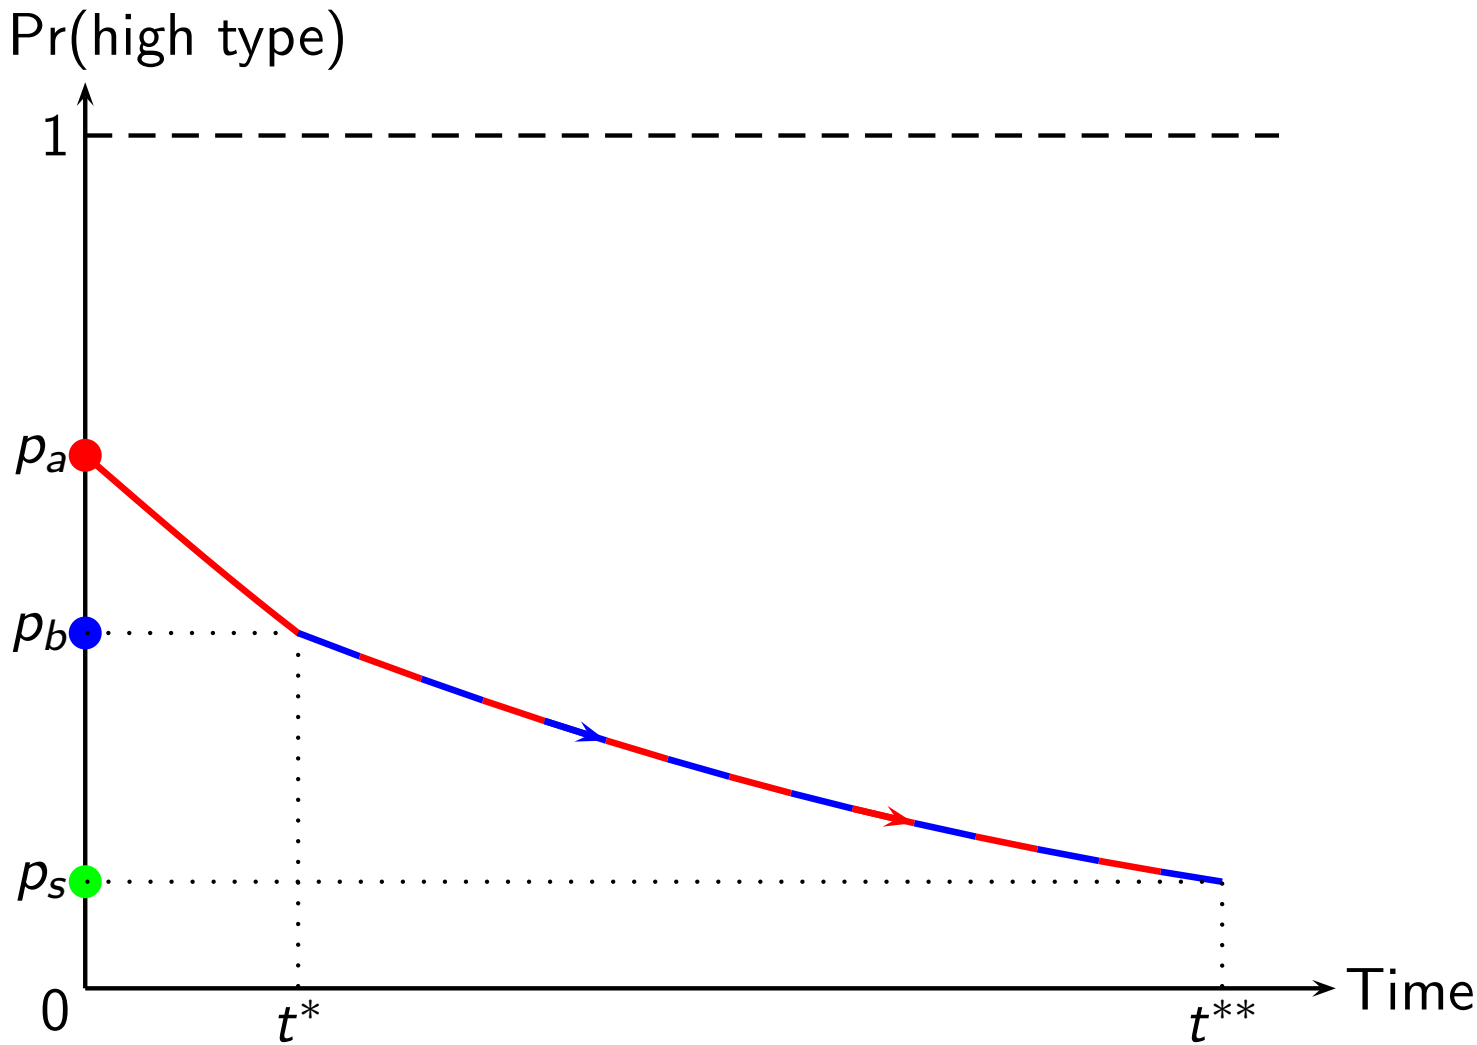

Breakthrough learning

Employer's optimal strategy:

- allocate the task to worker a over $[0, t^*)$
- mix equally over $[t^*, t^{**})$ if no breakthrough
- switch to safe arm at t^{**} if no breakthrough

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

Self-correcting under breakthrough learning

Proposition 1: *As $p_b \uparrow p_a$, worker b's expected payoff converges to worker a's expected payoff.*

- Task is assigned to worker a exclusively during $[0, t^*)$
- As $p_b \uparrow p_a$, $t^* \to 0$

Breakdown learning

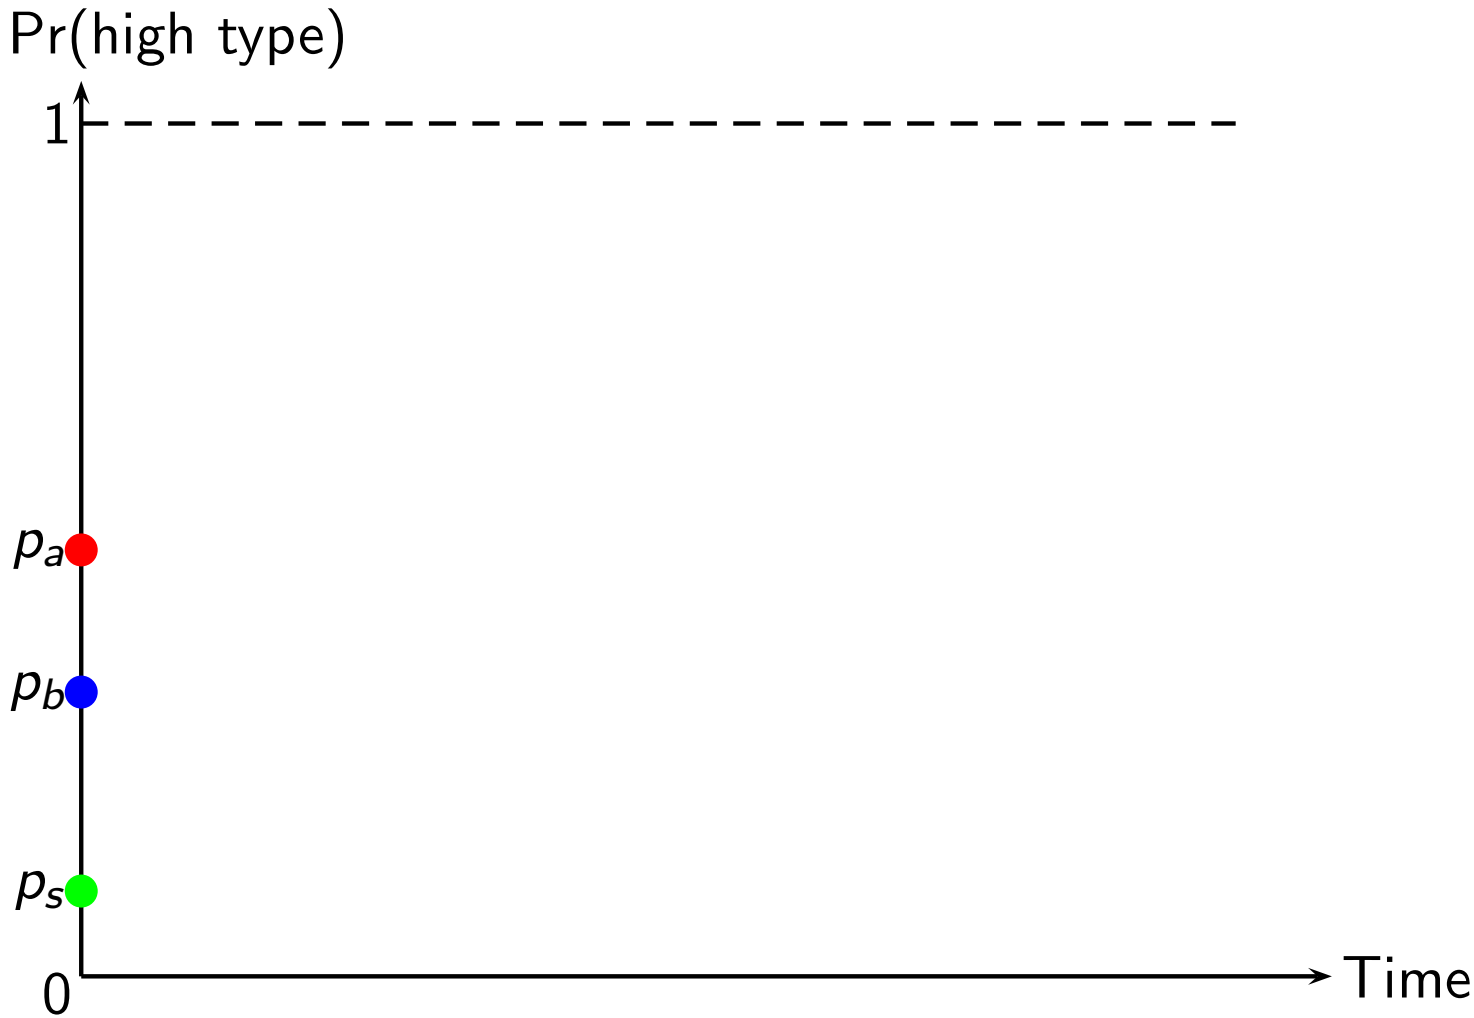

Breakdown learning

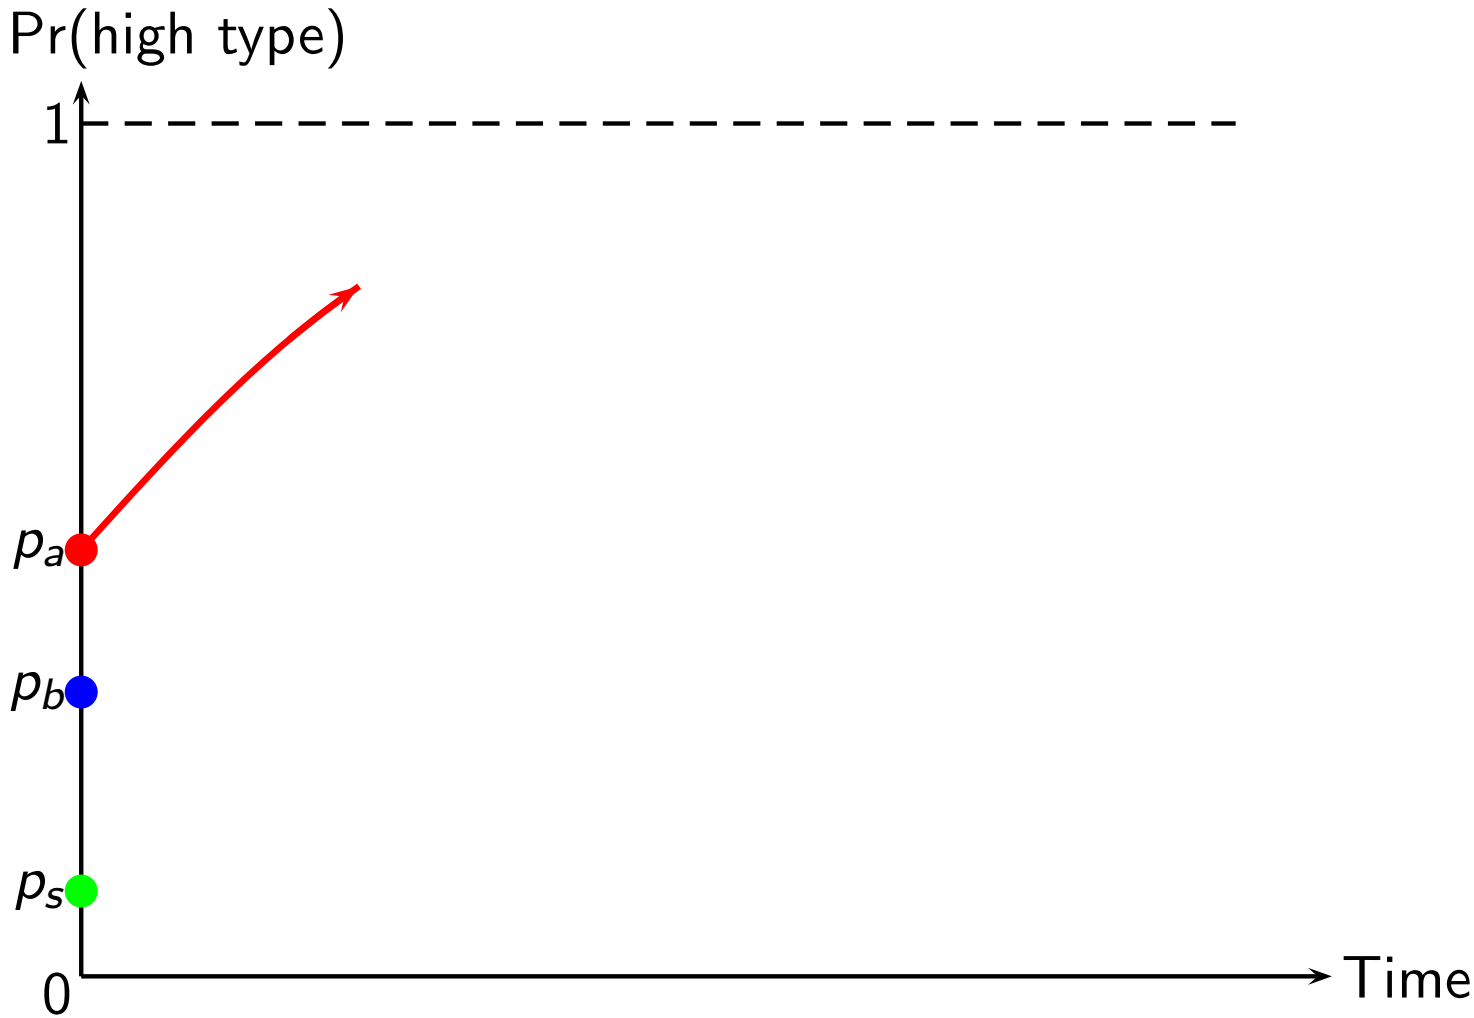

Breakdown learning

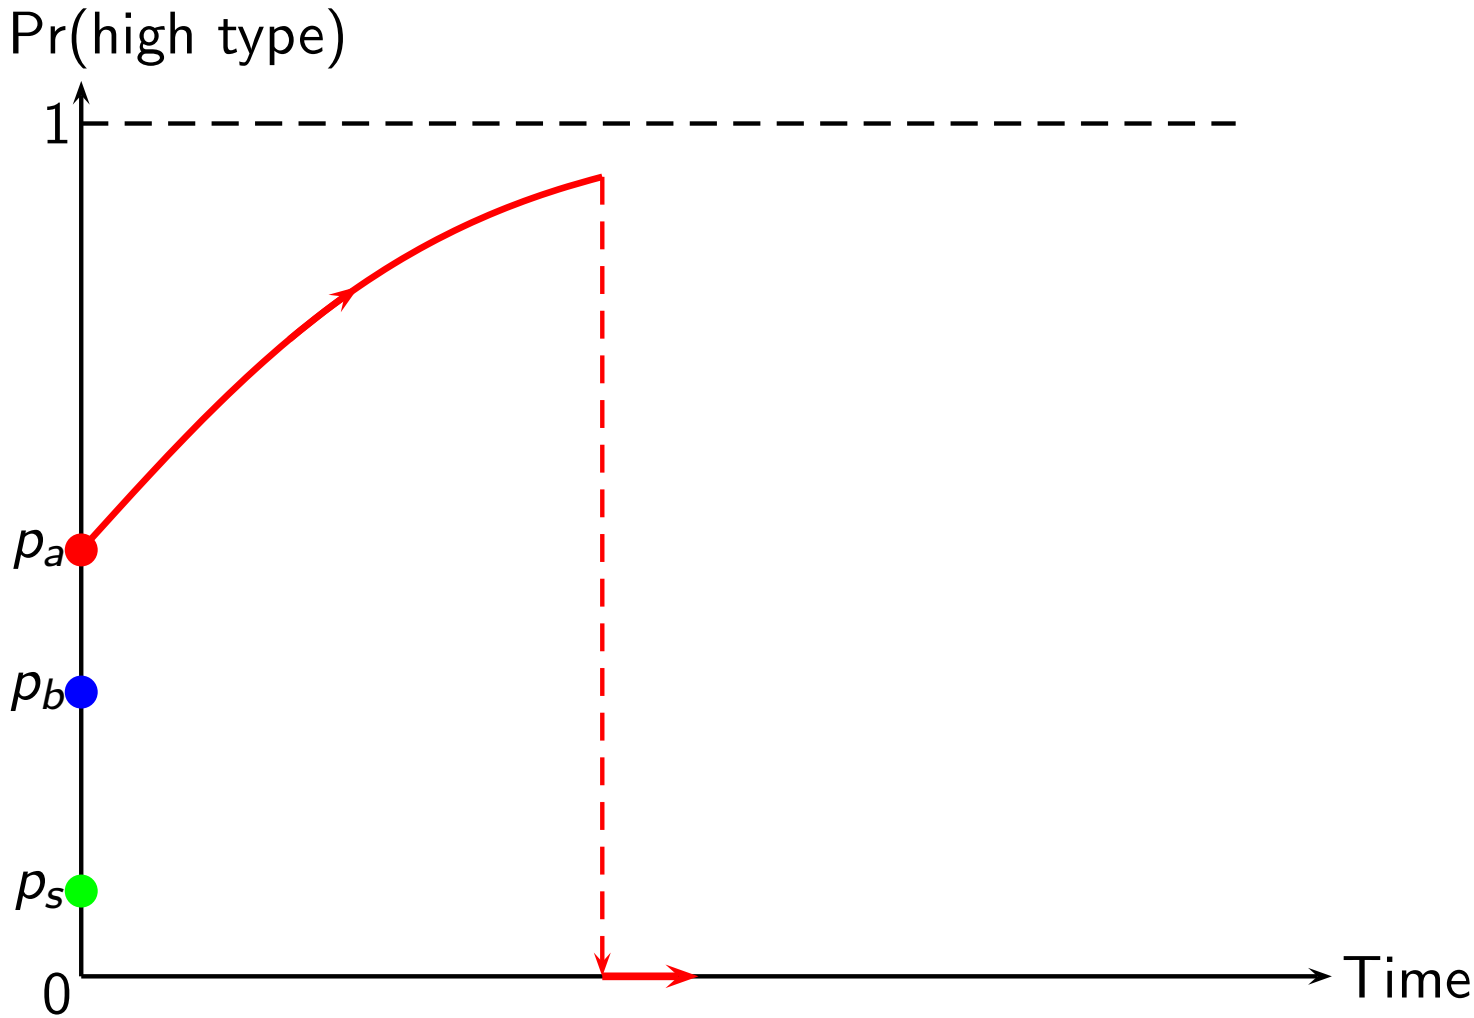

Breakdown learning

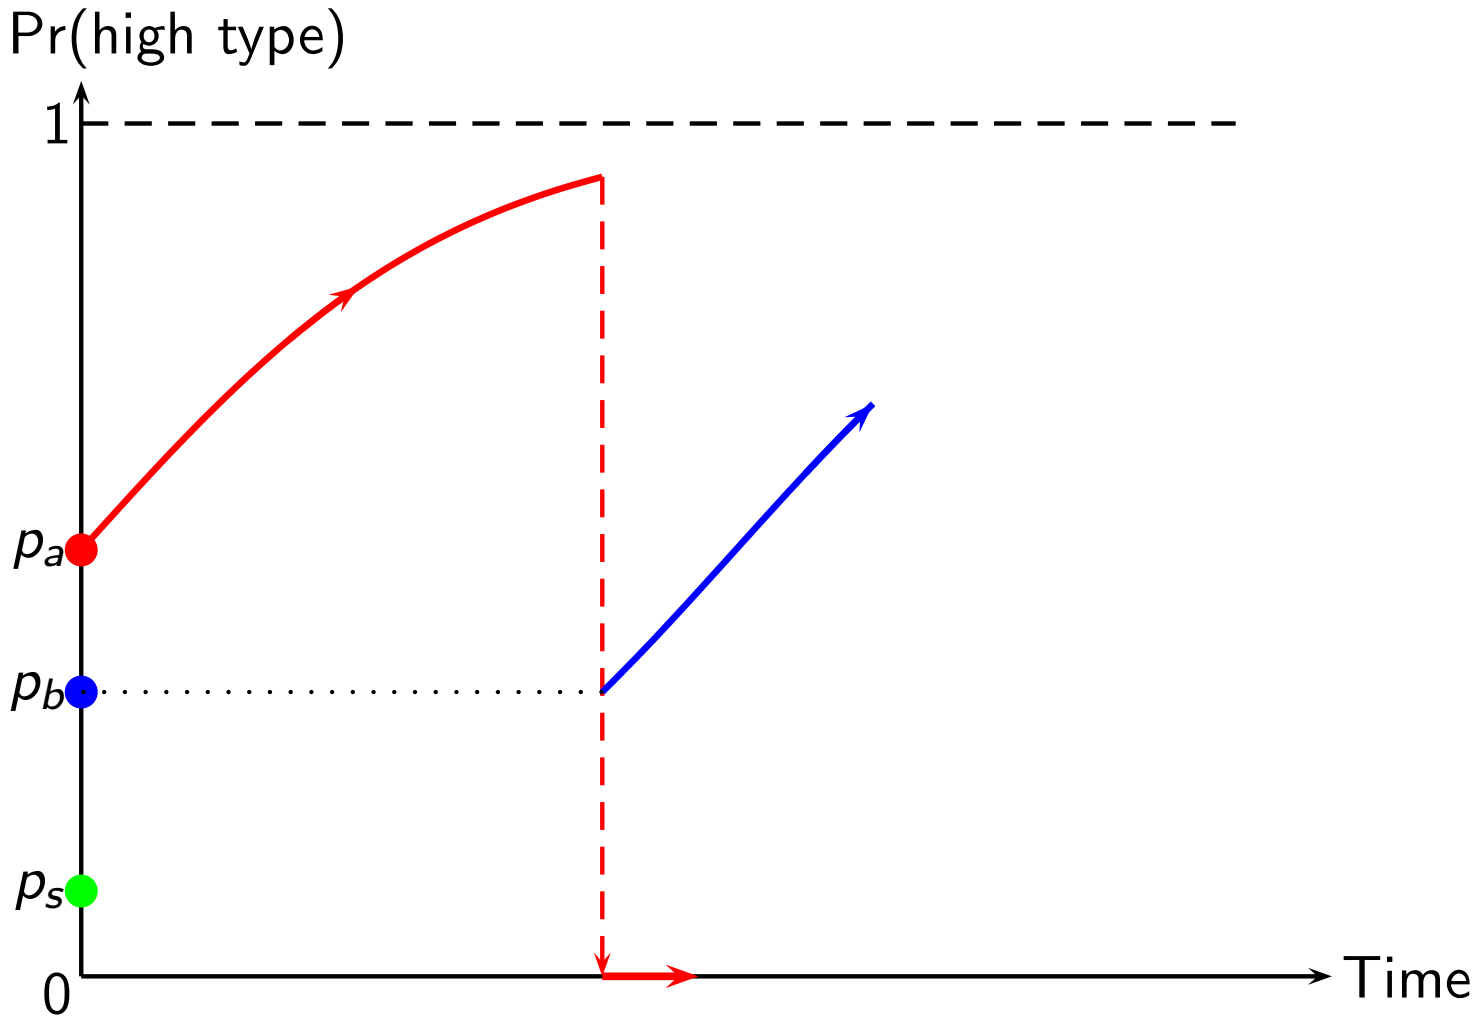

Breakdown learning

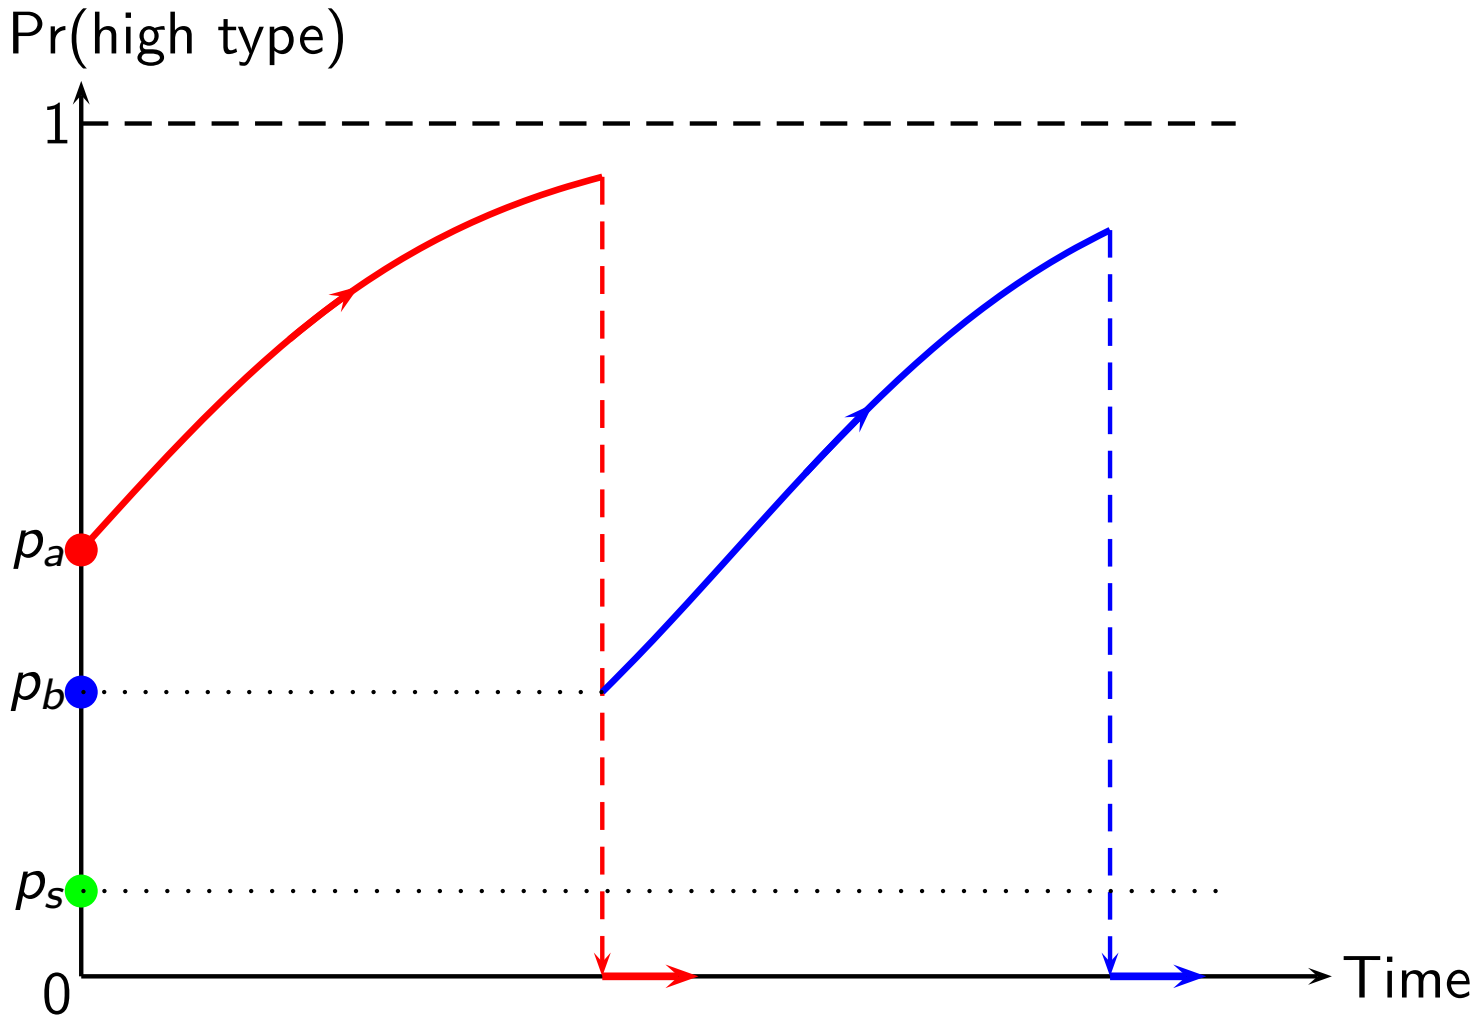

Breakdown learning

Employer's optimal strategy:

- allocate the task to worker *a* for as long as no breakdown occurs
- switch to worker *b* if/when worker *a* generates a breakdown
- switch to safe arm if both workers generate breakdowns

Spiraling under breakdown learning

Proposition 2: *As $p_b \uparrow p_a$, the ratio of worker b's expected payoff to worker a's expected payoff approaches*

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

- Task is assigned to worker a until he generates a breakdown
- Worker a's payoff

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

- Worker b's payoff

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

Alternative assumptions

Small initial difference:

- has no long-term impact under breakthrough learning
- grants disproportionate advantage to worker *a* under breakdown learning
  - even as $\lambda_\ell \to \infty$

Robust to:

- misspecified beliefs: $p_b = p_a$ but $\tilde{p}_b < \tilde{p}_a$ (Bohren, Imas and Rosenberg, 2019)
- inconclusive breakthroughs/breakdowns
- large labor markets: many employers and workers
- flexible wages

Roadmap

- Baseline model
- Self-correcting vs spiraling
- Large labor markets
- Flexible wages
- Investment in productivity

Large labor markets

- Unit mass of employers (tasks), α mass *a*-workers, β mass *b*-workers
- Relative task scarcity: more workers than tasks (for the talk: $\alpha > 1$)
- Frictionless matching

Large labor markets

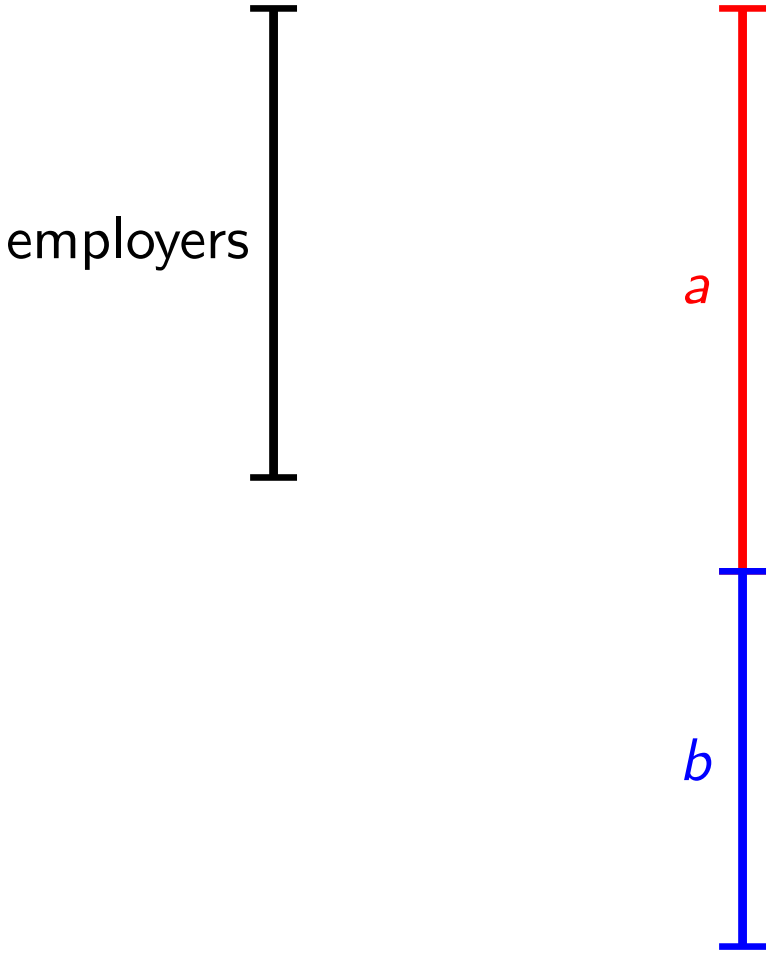

Breakthrough learning: broad hiring

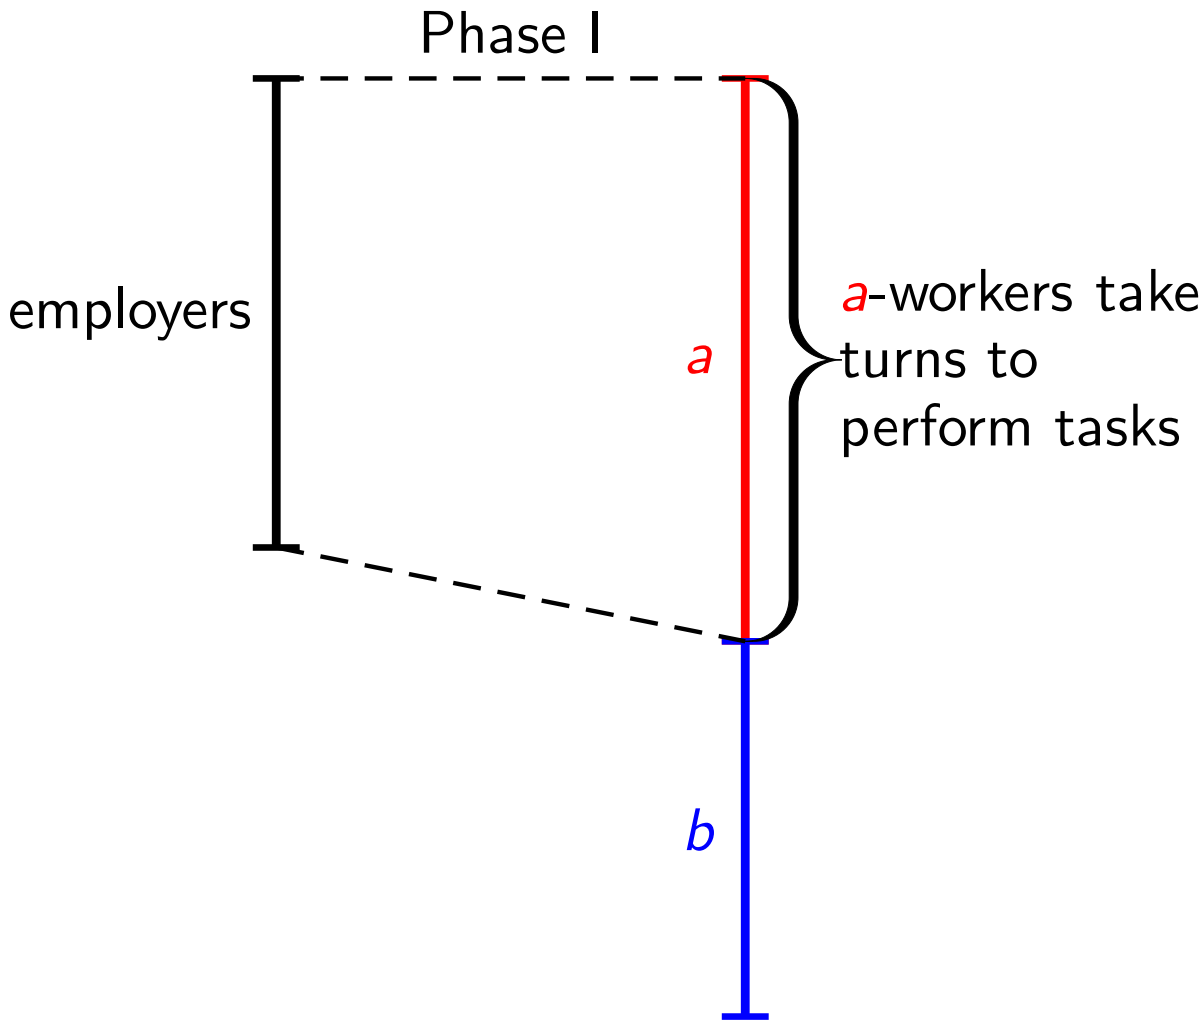

Breakthrough learning: broad hiring

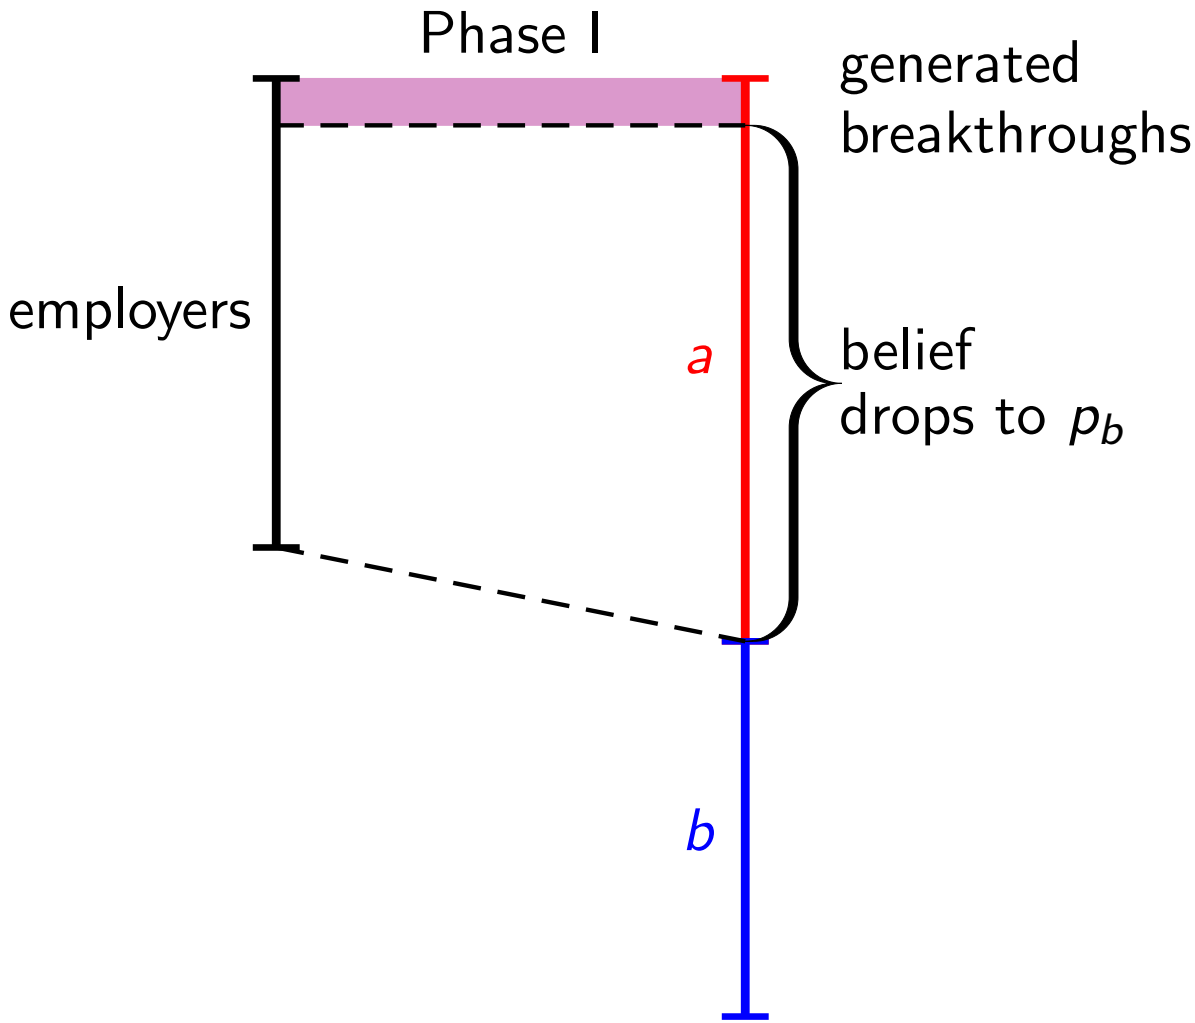

Breakthrough learning: broad hiring

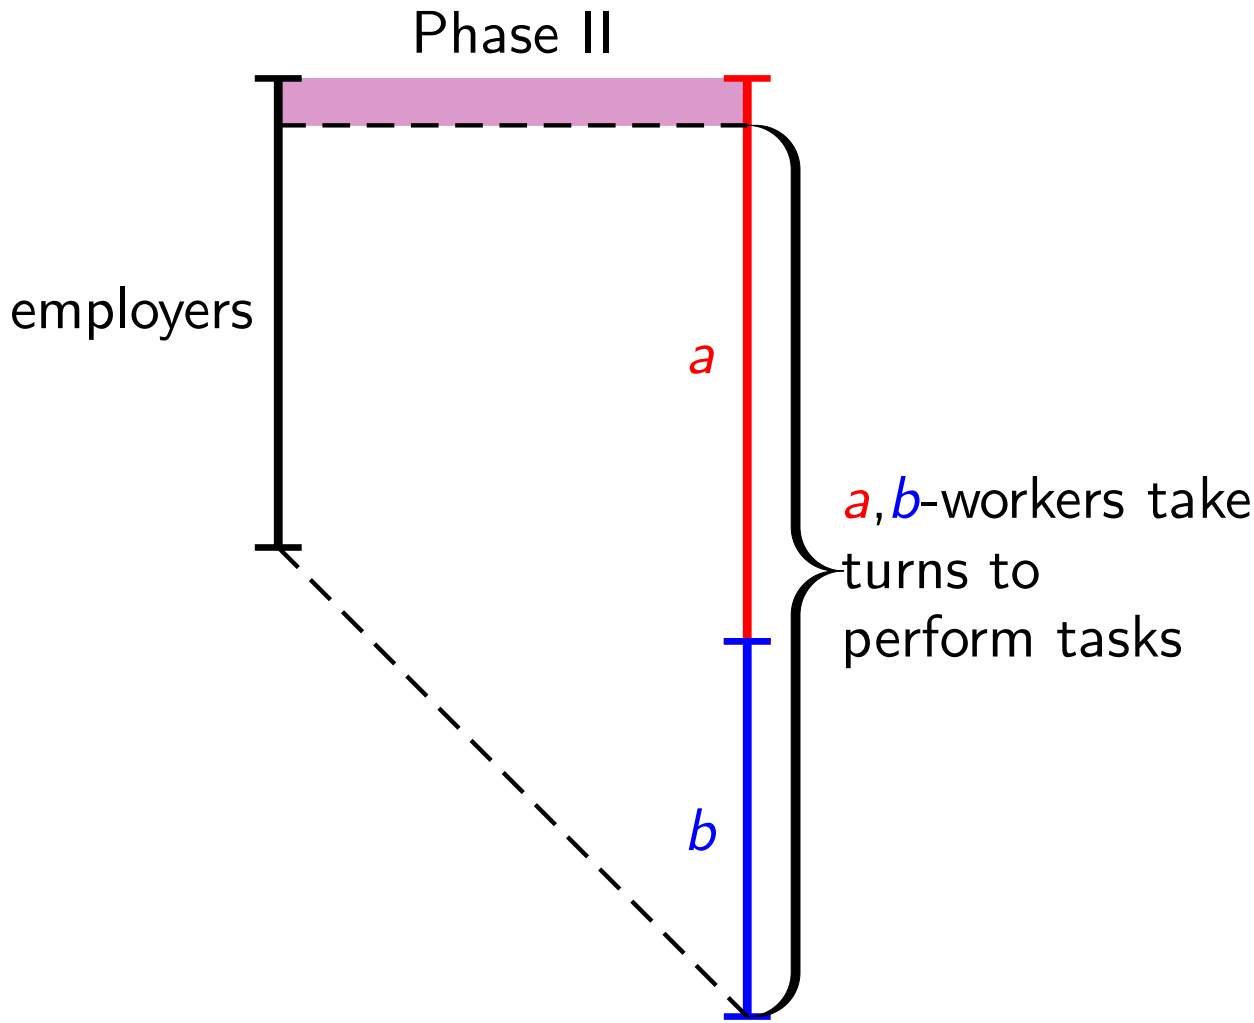

Breakthrough learning: broad hiring

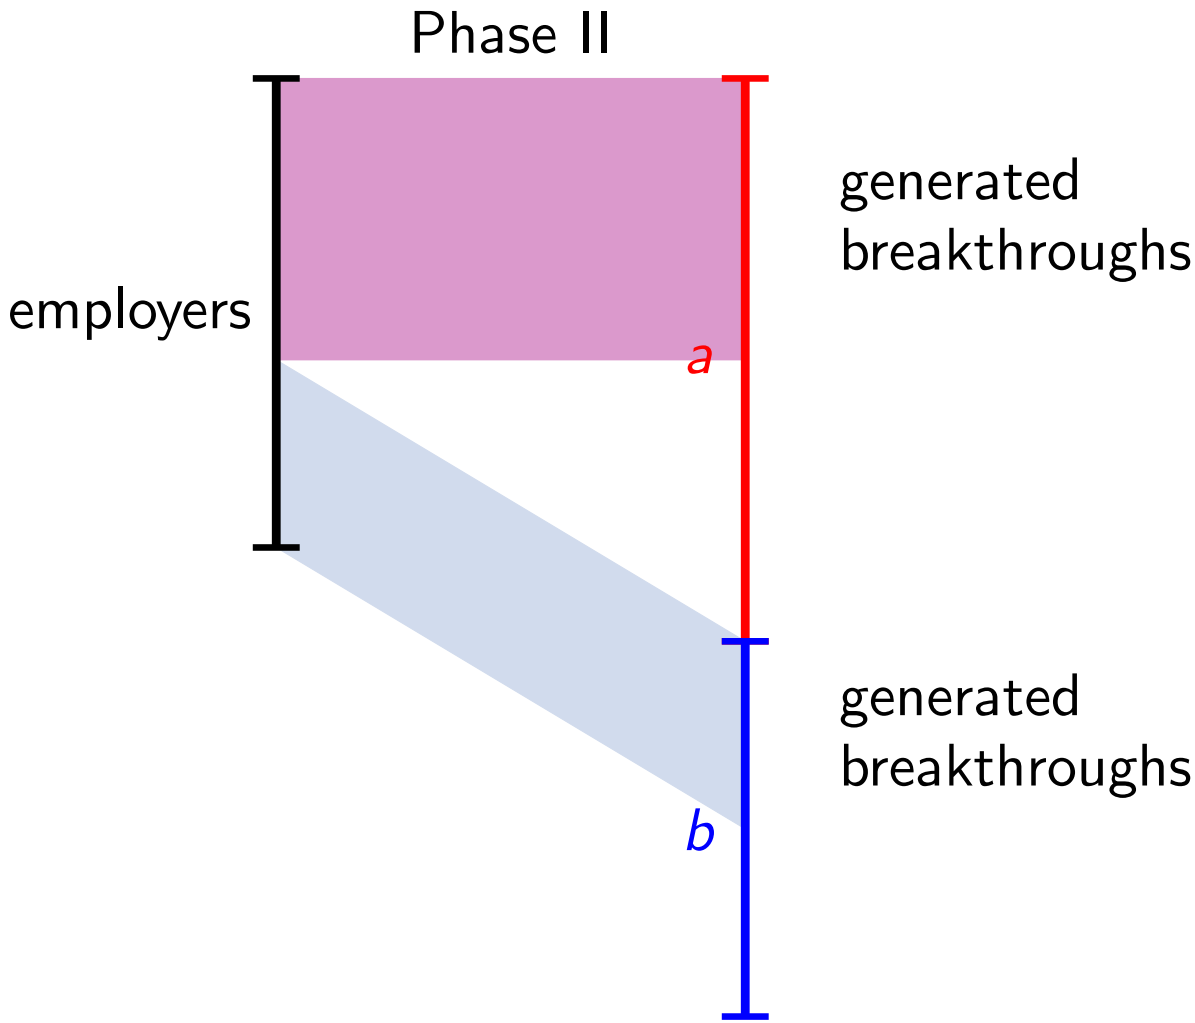

Breakthrough learning: broad hiring

'For a star job, the costs of a hiring error are small relative to the upside potential from finding an exceptional individual. Therefore, the organization will wish to sample widely among many employees, looking for the one pearl among the pebbles.'

Baron and Kreps (1999)

As $p_b \uparrow p_a$, a-workers and b-workers have the same expected payoff

Breakdown learning: narrow hiring

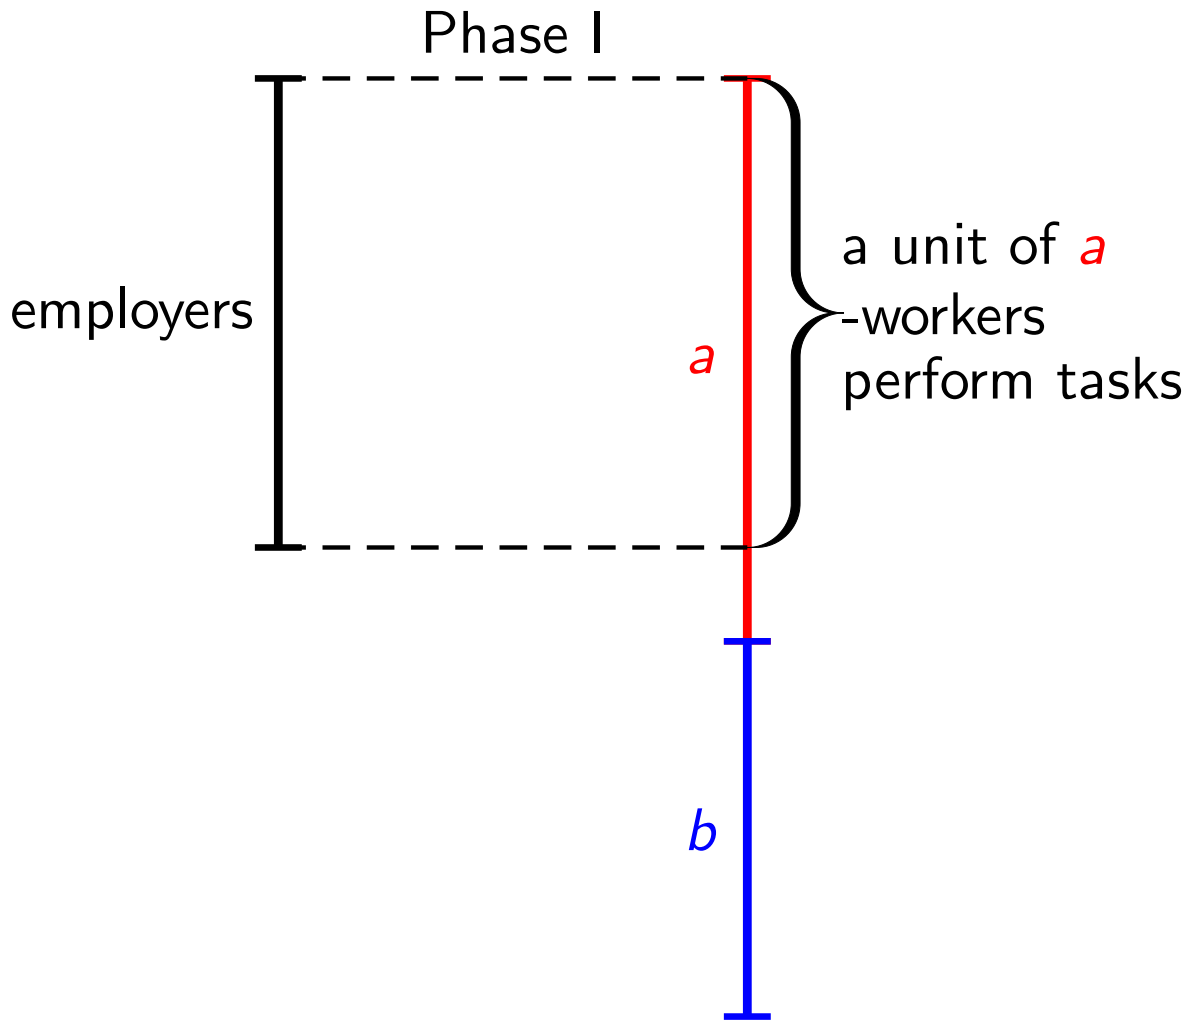

Breakdown learning: narrow hiring

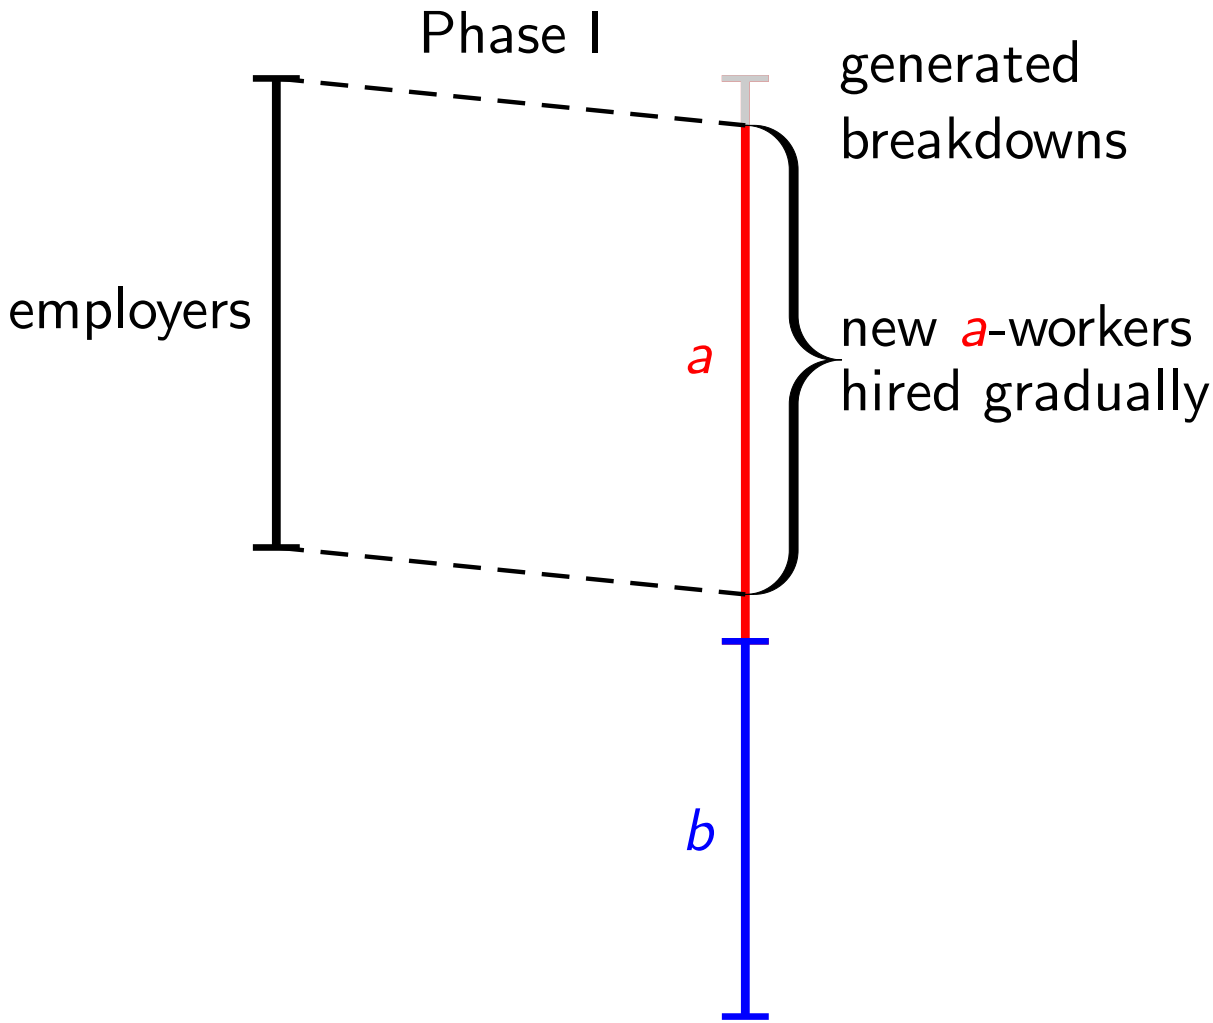

Breakdown learning: narrow hiring

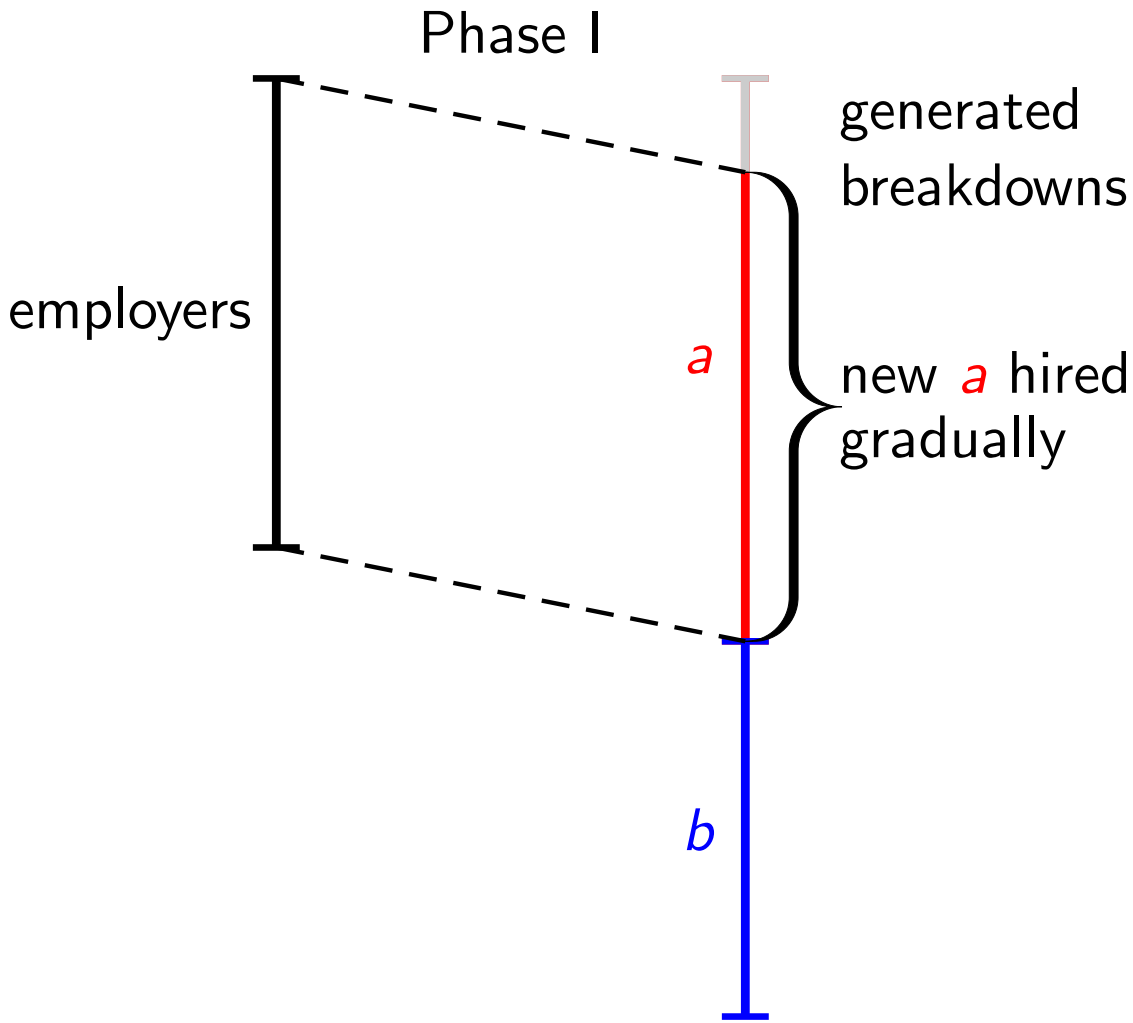

Breakdown learning: narrow hiring

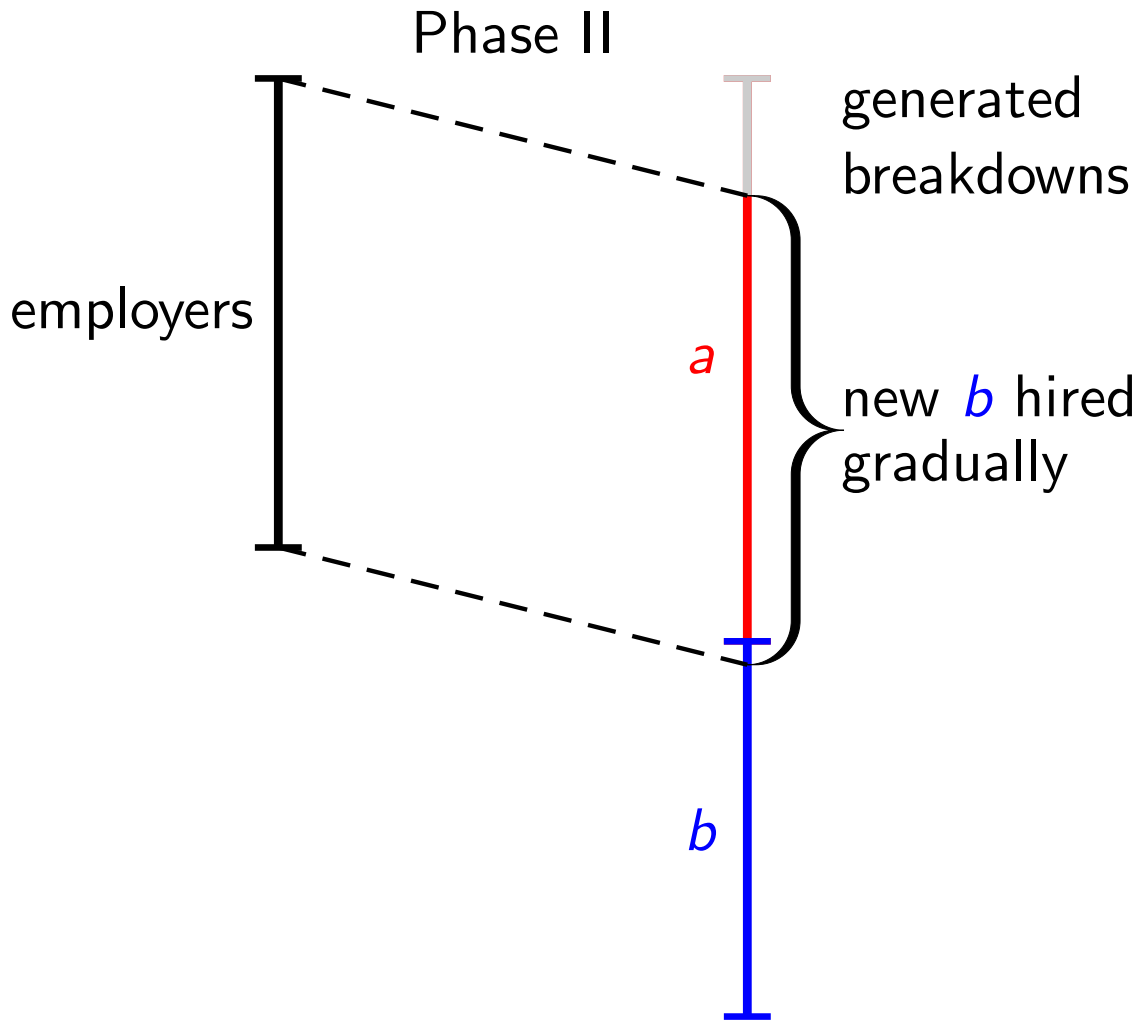

Breakdown learning: narrow hiring

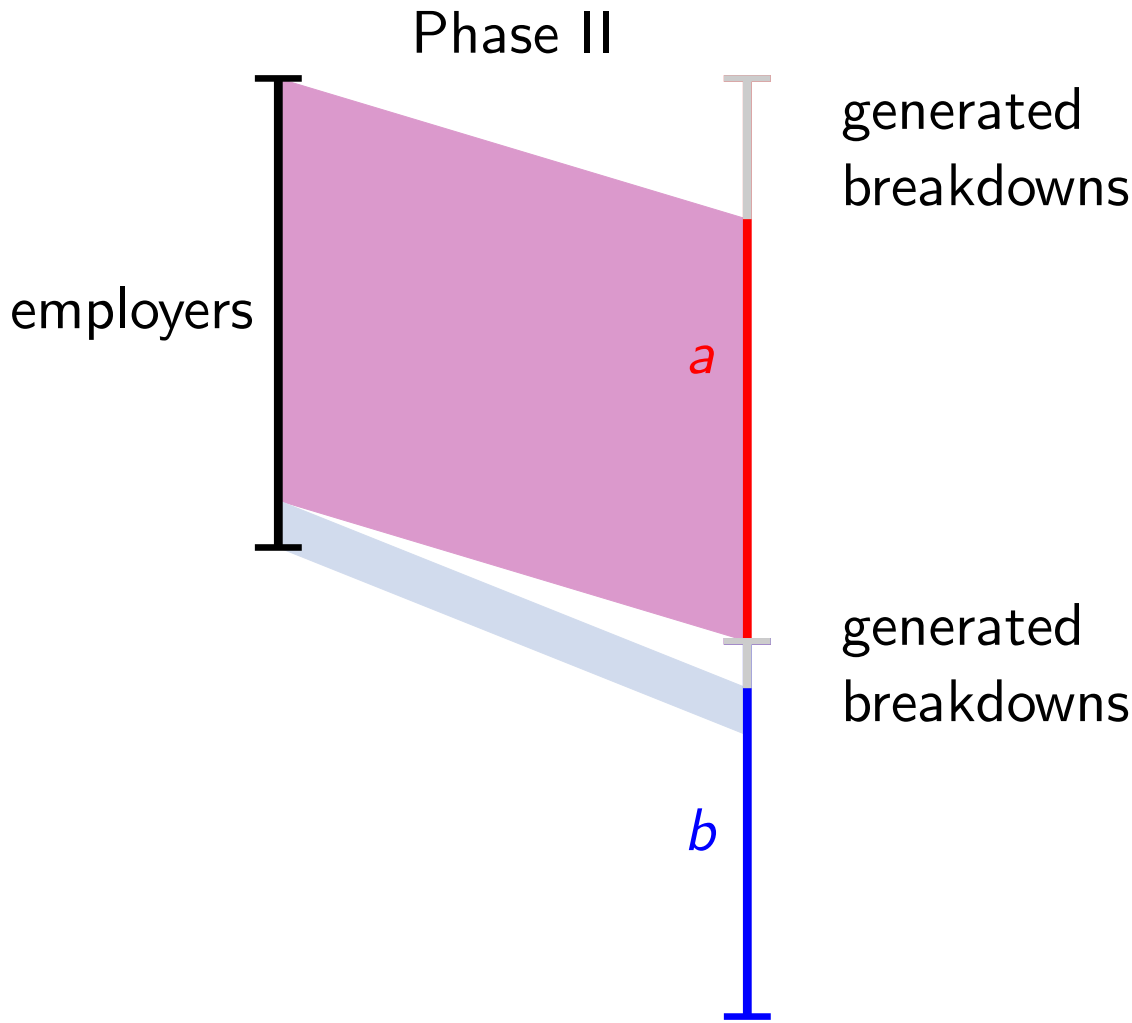

Breakdown learning: narrow hiring

- This process is the opposite of "sampling widely" under breakthrough learning
- b-workers are hired only after sufficiently many a-workers failed
- Even if $p_b \uparrow p_a$, delay in employment for b-workers does not vanish, so b-workers expect a smaller payoff than a-workers do

Breakdown learning: larger labor supply, greater inequality

Proposition 3: *As $p_b \uparrow p_a$, the ratio of the expected payoff of a b-worker to that of an a-worker decreases in both α and β.*

- $\beta \Uparrow$: intensifies competition among *b*-workers without affecting *a*-workers
- $\alpha \Uparrow$: intensifies competition among *a*-workers, and increases the delay for *b*-workers

  adding one *a*-worker uniformly delays every *b*-worker's employment

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

Roadmap

- Baseline model
- Self-correcting vs spiraling
- Large labor markets
- Flexible wages
- Investment in productivity

Flexible wages: framework

- Introduce flexible wages to the large market
- A stage-game outcome specifies
  - how workers are matched to employers
  - a nonnegative wage for each matched pair
- We characterize the stable stage-game matching (Shapley and Shubik, 1971), which is essentially unique
- Prescribing the stable stage-game matching after each history is dynamically stable (Ali and Liu, 2020)

Flexible wages: solution

- There is a history-dependent marginal productivity p^M
- A worker is matched iff his expected productivity p exceeds p^M
  - his wage is $\left(p - p^M\right) v$
- An unmatched worker receives 0

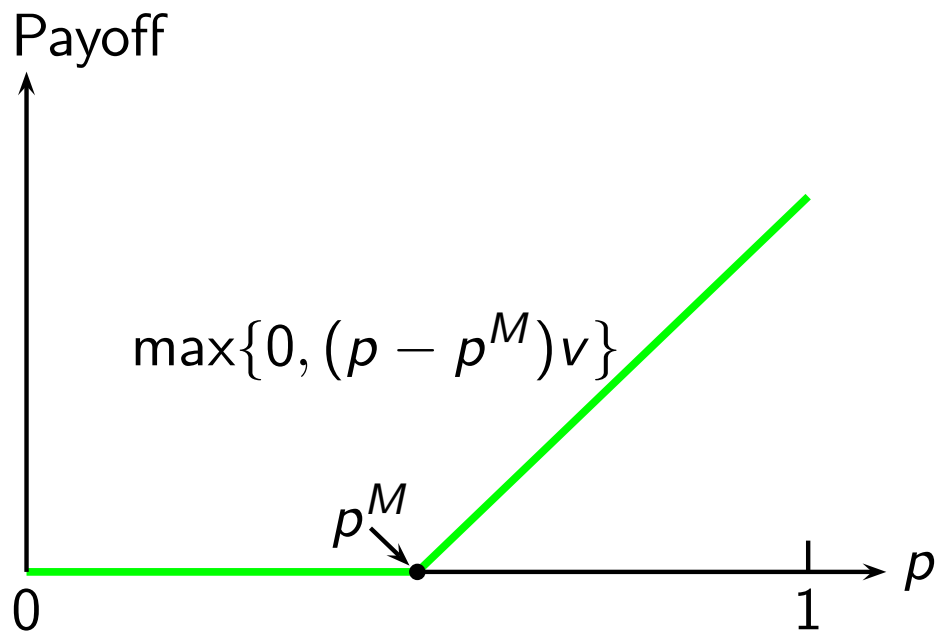

Flexible wages do not fix spiraling under breakdowns

- More learning about a worker's type $\Longrightarrow$ higher expected payoff
- Delay in employment for b-workers does not vanish as $p_b \uparrow p_a$
- More is learnt about a-workers than b-workers

Flexible wages do not fix spiraling under breakdowns

Two-period intuition: $\alpha = \beta = 1$

Flexible wages do not fix spiraling under breakdowns

Two-period intuition: $\alpha = \beta = 1$

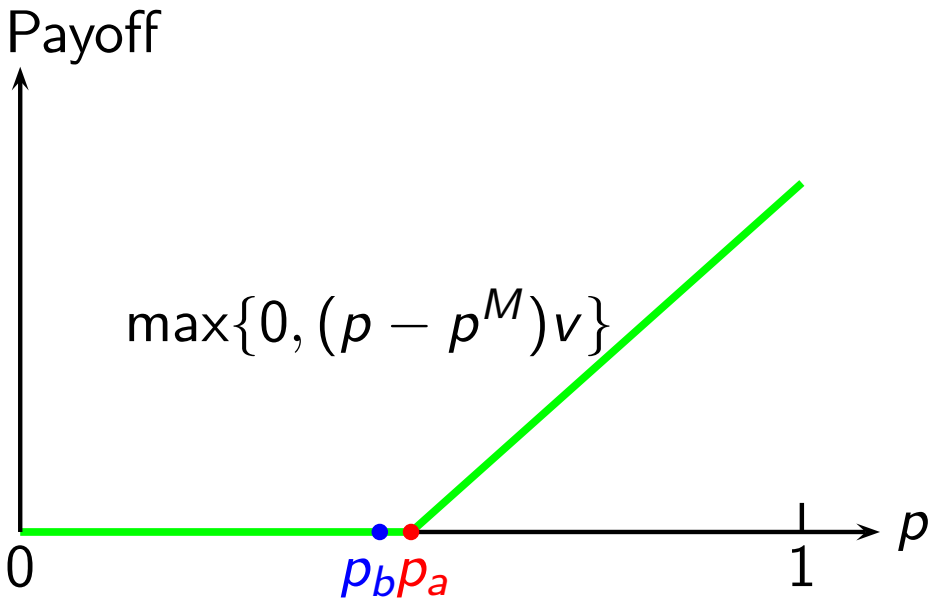

First period

Flexible wages do not fix spiraling under breakdowns

Two-period intuition: $\alpha = \beta = 1$

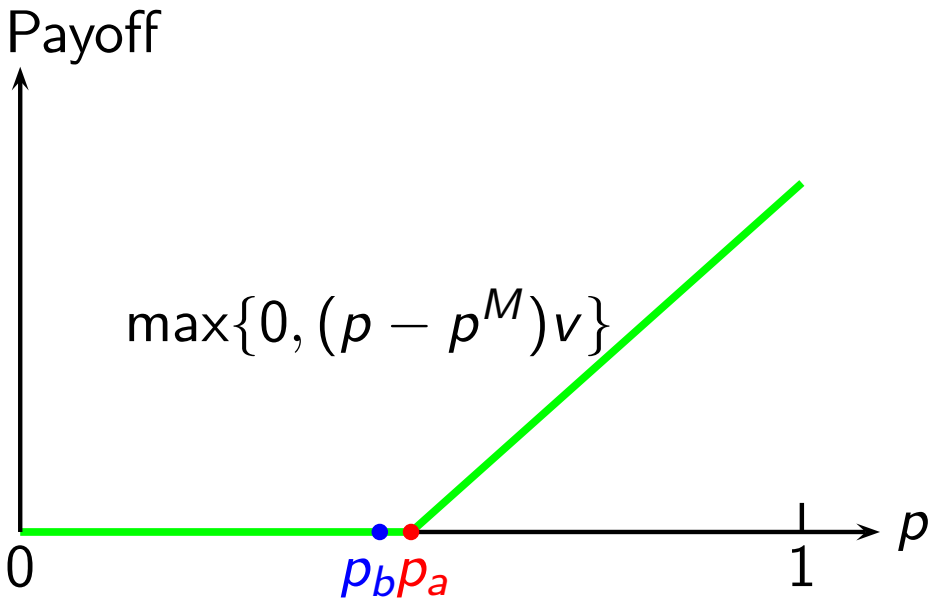

First period

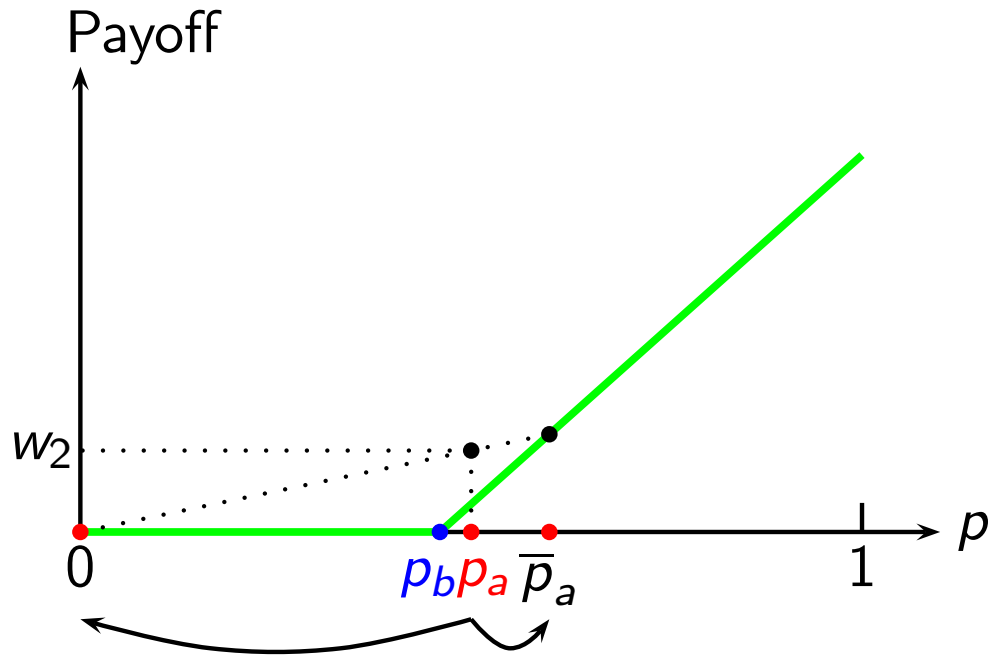

Second period

Persistent gaps in average wages and payoffs even if $p_b \uparrow p_a$

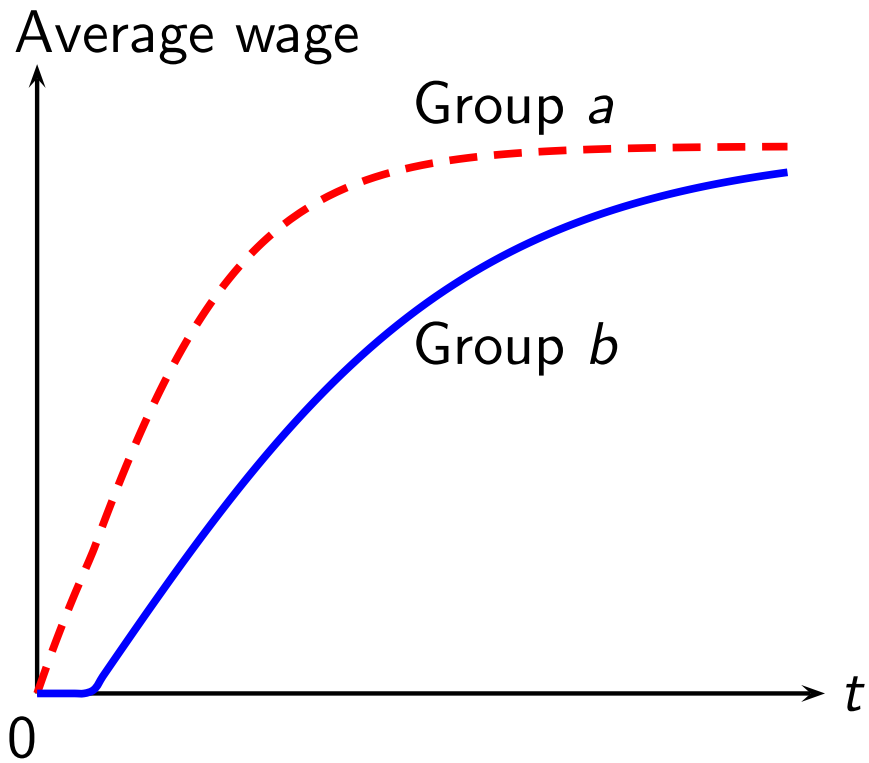

conclusion

Persistent gaps in average wages and payoffs even if $p_b \uparrow p_a$

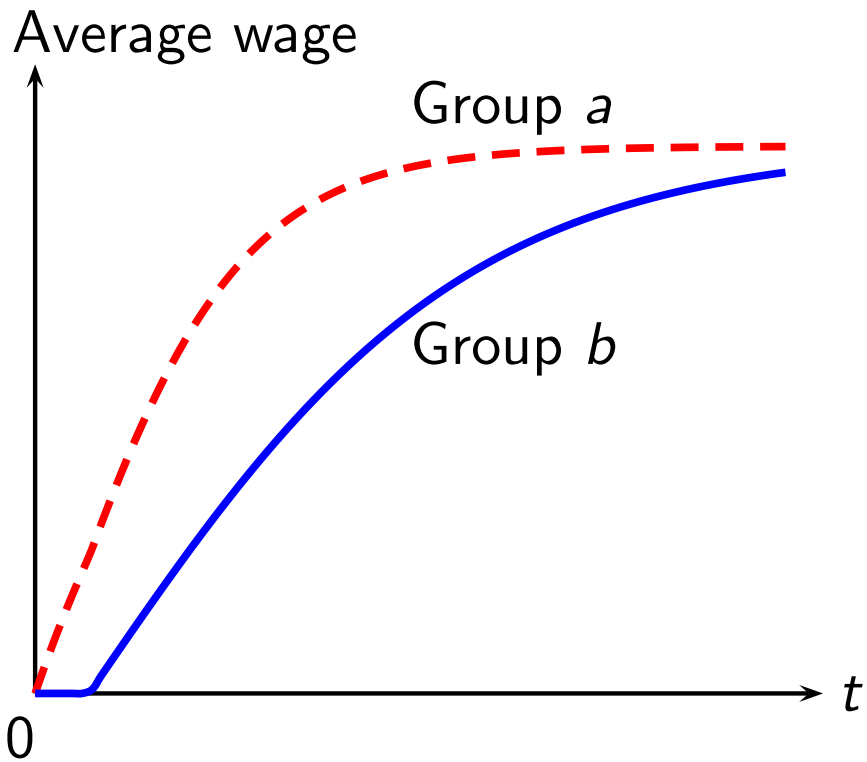

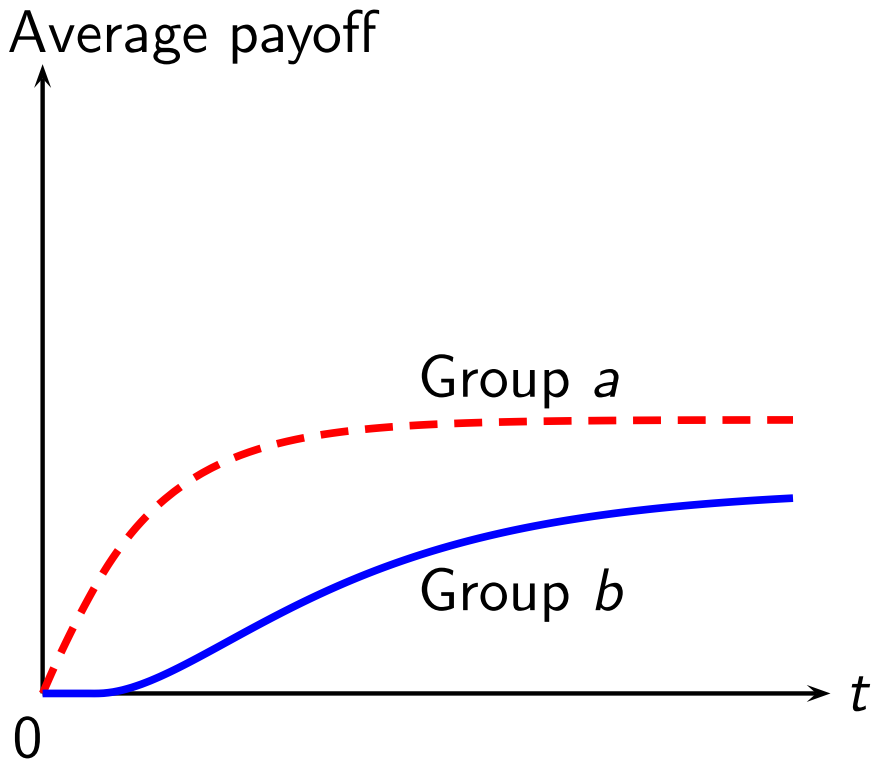

conclusion

Roadmap

- Baseline model
- Self-correcting vs spiraling
- Large labor markets
- Flexible wages
- Investment in productivity

Investment in productivity

- Before $t = 0$, each low-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 and the investment decision are private information to the worker
- F is the same for both workers

Investment equilibrium

- Let (q_a, q_b) denote the employer's post-investment belief
- The employer chooses an optimal allocation strategy given (q_a, q_b)
- Worker i' benefit from investment is:

$$B_i(q_a, q_b) = U_i(h; q_a, q_b) - U_i(\ell; q_a, q_b),$$

  where $U_i(\theta_i; q_a, q_b)$ is worker i's payoff given (q_a, q_b) and type θ_i
- Worker i invests if and only if his cost is below the benefit $B_i(q_a, q_b)$
- An equilibrium is a pair of beliefs (q_a, q_b) satisfying:

$$q_i = p_i + (1 - p_i)F(B_i(q_a, q_b)), \ \forall i \in \{a, b\}$$

What is common between the two learning 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
- We compare the equilibrium sets across two learning environments

Preview of key results

Result 1: equilibrium lifetime payoffs for workers

- Investment does not disturb the self-correcting property of breakthrough learning
- Investment exacerbates spiraling under breakdown learning: it makes the workers' payoffs more unequal than without investment
  - breakdown learning is "cancer," investment is "complication"

Result 2: investment behavior

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

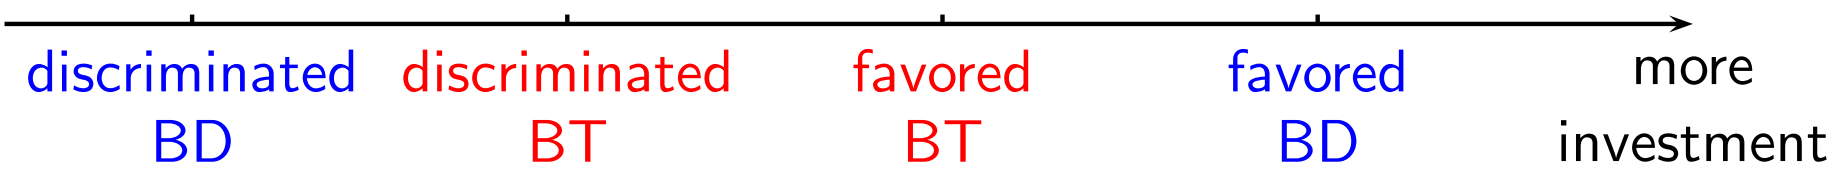

Self-correcting property survives

Proposition 4: *As $p_b \uparrow p_a$, there exists an equilibrium in which the two workers' expected payoffs as well as their post-investment beliefs converge.*

Proof idea:

- When $p_a = p_b$, there exists an equilibrium in which $q_a = q_b$
- The benefit from investment $B_i(q_a, q_b)$ is continuously differentiable.
- By implicit function theorem, when $p_b \uparrow p_a$, there exists an equilibrium in which $q_b \to q_a$

Exacerbated spiraling under breakdowns

Proposition 5: *As $p_b \uparrow p_a$, in any equilibrium, the ratio of discriminated worker's payoff to favored worker's payoff is at most*

$$(1 - q_i)\frac{\lambda_\ell}{\lambda_\ell + r} < 1,$$

where q_i is favored worker's belief. This ratio is strictly below the ratio in the no-investment benchmark.

Proof idea:

- The payoff ratio is pinned down by how likely it is that the favored worker has a high type
- Favored worker is more likely to be a high type with investment than without

Investment polarization under breakdowns

Proposition 6: *There exists $\bar{\lambda} > 0$ such that for any $\lambda_h, \lambda_\ell \geqslant \bar{\lambda}$ and in any pair of equilibria, one from each environment, favored worker invests strictly more under breakdowns than under breakthroughs, while discriminated worker invests strictly less.*

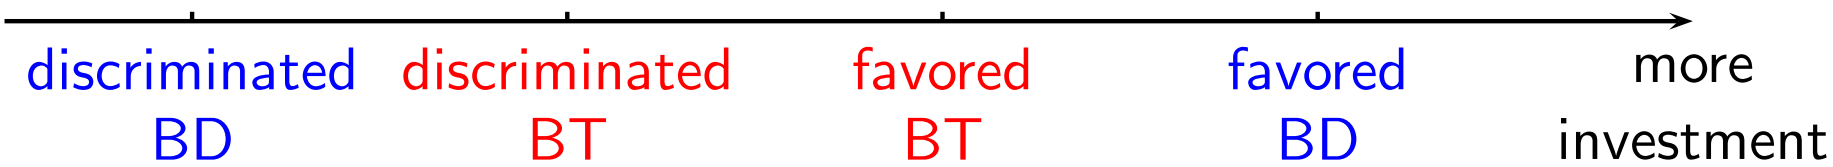

Investment polarization under breakdowns

Proof idea:

Favored worker invests more under breakdown learning:

- Under breakdown learning, the benefit from investment for favored worker is close to one

  A high type gets 1, and a low type gets 0
- Under breakthrough learning, the benefit is < 1

  A high type gets < 1, and a low type gets 0

Discriminated worker invests less under breakdowns:

- Breakdown learning already disfavors the second worker to be hired
- Favored worker invests strictly more under breakdown learning

Employer's preferred learning environment

Corollary: *There exists $\bar{\lambda} > 0$ such that for any $\lambda_h, \lambda_\ell \geqslant \bar{\lambda}$, the employer's payoff is strictly higher under breakdown learning than under breakthrough learning.*

Proof idea:

- Under breakdown learning, favored worker invests almost for sure
- So employer is guaranteed to hire a high type

Tradeoff: *efficiency* for employers vs. *equality* between the workers

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

Final thoughts

- More empirical work needed on the contrast between breakthrough and breakdown learning

  'Thus, while one challenge is to explain earnings differentials between black and white men, there is an even greater challenge, which is to explain the simultaneous existence of wage differentials among relatively low-skill male workers and their possible absence among high-skill male workers.' Lang and Lehmann (2012)

Thank you!

Interpreting breakthrough vs breakdown learning

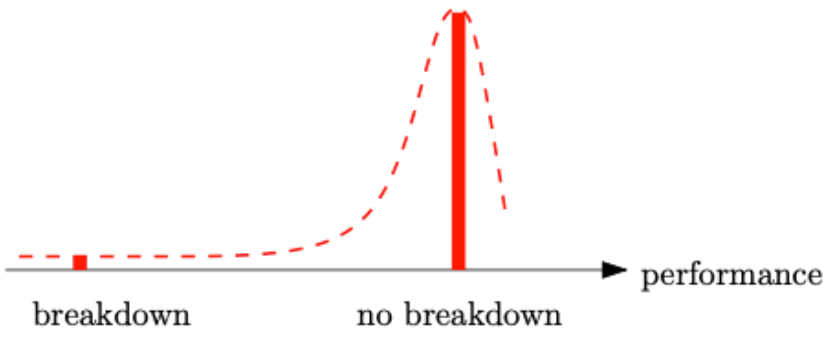

(a) Guardian jobs

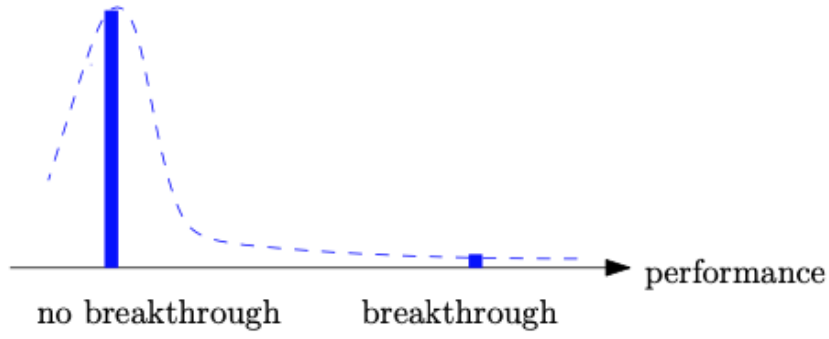

(b) Star jobs

Distribution of performance signals (fig. 2-2 in Baron and Kreps (1999))