The Distributional and Welfare Implications of Personalized Wages

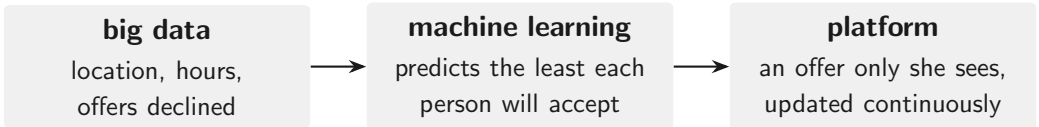
Robert Shimer
University of Chicago

XXIX Annual Conference of the Central Bank of Chile
November 20, 2026

A New Worry

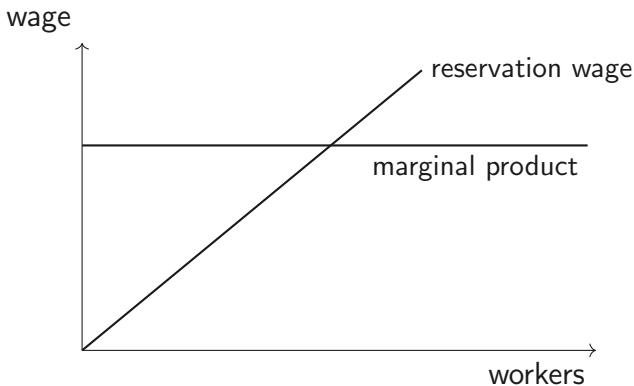
- firms pay different workers different wages for similar work
  - "granular data on location, individual behavior, demand, supply, or other factors"
  - "algorithmic wage discrimination" (Dubal 2023)
- October 2026: FTC orders to Uber, Lyft, DoorDash, and Instacart
  - how do they use workers' personal data to set pay?
  - "such as their phone usage or acceptance history for past offers"
- same worry about personalized/surveillance pricing (CEA 2015; FTC 2025, 2026)
  - companion paper: *The General-Equilibrium Incidence of Personalized Pricing*
- the fear: a firm will always hold you to your **reservation wage**
  - income shifts to those who harness the information

Why Now? The Missing Ingredient Was Information



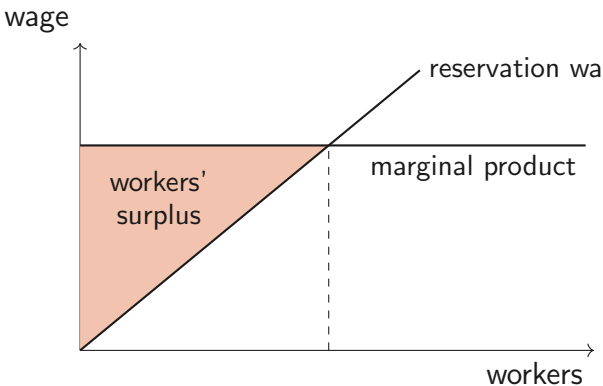
- paying each person the least she will accept is an old idea (Pigou 1920)
  - it stayed in the textbook: no employer knew each worker's reservation wage
- suppose preferences, technology, and the number of competitors are constant
  - only information is changing
- **what does that do to wages, profits, and workers' well-being?**

A Labor Market



workers ranked by the lowest wage each will accept; every worker produces the same output

The Neoclassical Paradigm: Everyone Takes the Wage as Given



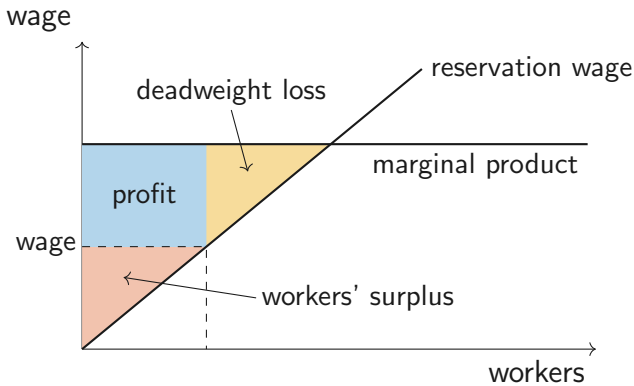
workers employer lost

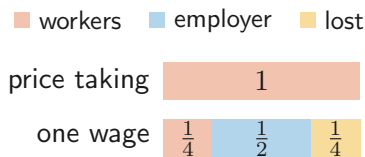
price taking 1

shares of the surplus from employment

a competitive market: one wage, the same for everyone, taken as given

Market Power: Still One Wage for Everyone

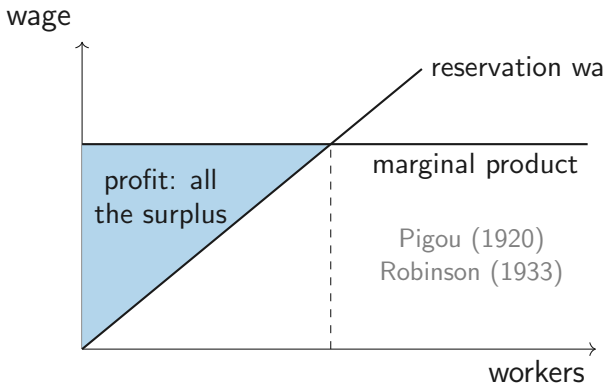


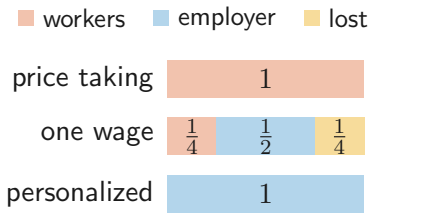


shares of the surplus from employment

a single employer sets the wage, but it is still one wage for every worker
all but the marginal worker are paid more than the least they would accept

Personalized Wages: The Textbook





shares of the surplus from employment

a wage for each person, not for the market
efficient employment, and the employer takes everything

Search Theory Has Personalized Wages, Not Personalized Tastes

	a worker's wage depends on	matches differ in
bargaining	her outside option	productivity
wage posting	firms' posted wage	productivity
sequential auctions	her best competing offer	productivity
the personalization debate	the least she will accept	**her taste for each job**

Mortensen–Pissarides (1994); Burdett–Mortensen (1999); Postel-Vinay and Robin (2002)

- search theory gave up the law of one price long ago
- an employer that knows a worker's alternatives pays just enough to beat them
  - but her alternatives depend on productivity and luck, not on which jobs she likes
- the debate is about employers learning how much each worker wants each job

Two Questions

1. if firms set wages for people, not for markets, what is the right paradigm?
   - the same question for prices

2. in that paradigm, does income shift toward those who harness the information?
   - the textbook: yes, all of it

One Proposal for a Paradigm

- people differ in idiosyncratic **tastes**: for each job, for each good
- tastes give firms **market power** even when no firm observes them
  - posted wages: monopsony (Card et al. 2018, Berger et al. 2022)
  - uniform prices: monopolistic competition (Perloff—Salop 1985)
- **information** is a primitive, like preferences, endowments, and technology
  - what each firm observes about each person, or is allowed to use
- **equilibrium**: every firm is informed, and every other firm responds
  - the textbook holds everything else fixed

Does Income Shift to Those Who Harness the Information?

	prices (companion paper)	**wages** (this talk)
firms learn	her taste for each good	her taste for each job
the textbook	sellers take her surplus	employers pay her reservation wage

Does Income Shift to Those Who Harness the Information? No

	prices (companion paper)	**wages** (this talk)
firms learn	her taste for each good	her taste for each job
the textbook	sellers take her surplus	employers pay her reservation wage
the textbook holds fixed	every other price and fee	every other employer

Does Income Shift to Those Who Harness the Information? No

	prices (companion paper)	**wages** (this talk)
firms learn	her taste for each good	her taste for each job
the textbook	sellers take her surplus	employers pay her reservation wage
the textbook holds fixed	every other price and fee	every other employer
the benchmark	same allocation, same profits	same average wage, same profits
beyond it	the poor gain, the rich lose	more employment, same labor share

- the firm she likes best must still beat the runner-up, which will pay her full product
- information changes **who bears risk and who works**, not who gets the surplus
- the exception: a single employer (Postel-Vinay and Robin, 2002)

Outline

1. Model

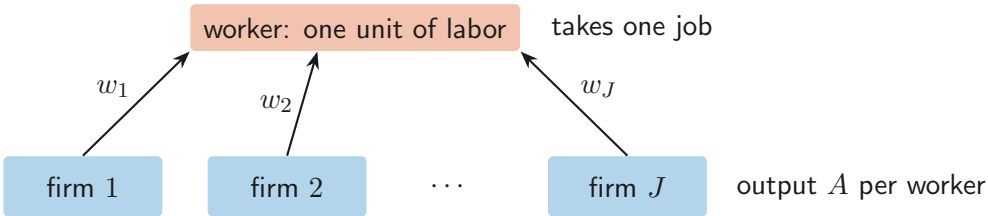
2. Wages and Profits

3. Welfare

4. Employment

5. Conclusion

Technology and Endowments



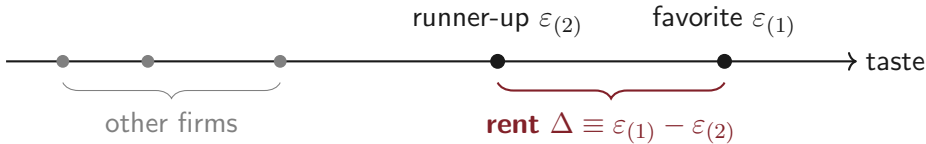
- a unit measure of workers, each with one indivisible unit of labor
  - she works at one firm; later, she may choose not to work
- $J \geq 2$ identical firms; a worker produces A at any of them
  - later: diminishing returns
- firm j earns $A - w_j$ on each worker it employs; its owners are not workers

Preferences

$$U(w, \varepsilon_j) = \frac{\big(w \, \exp(\varepsilon_j)\big)^{1-\gamma}}{1-\gamma}, \qquad \gamma \geq 0$$

- ε_j: her taste for firm j; $w \, \exp(\varepsilon_j)$: her **amenity-adjusted wage**
- she sees every offer and every taste, then takes the job with the highest $\log w_j + \varepsilon_j$
- γ: relative risk aversion over the amenity-adjusted wage
  - it governs how she ranks lotteries over jobs, not which job she takes
  - her job choices do not reveal it
- paper: $\phi\big(u(w) + \varepsilon_j\big)$, with u concave and ϕ increasing

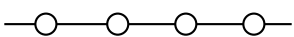
Tastes and the Worker's Rent



- ε_j: independent across firms and workers, exponential with mean and s.d. σ
  - the amenity $\exp(\varepsilon_j)$ has a Pareto distribution
- **rent**: how much more she likes her favorite firm than the runner-up
- memoryless: Δ is exponential with mean σ for any number of firms

Information and Equilibrium

posted wages



a firm observes none of her tastes

or: a ban on personalized pay

bidding



a firm observes her taste for its own job

or: no conditioning pay on her other options

personalized wages



every firm observes all of her tastes

no restriction

- same preferences, endowments, and technology: only information differs
- firms make wage offers simultaneously, then the worker chooses: Nash equilibrium
- price competition (Perloff–Salop 1985), a first-price auction, a second-price auction

Outline

1. Model

2. Wages and Profits

3. Welfare

4. Employment

5. Conclusion

Posted Wages: The Marginal Worker

- firm j maximizes $(A - w_j)\, P_j$, where P_j is the probability she takes its job

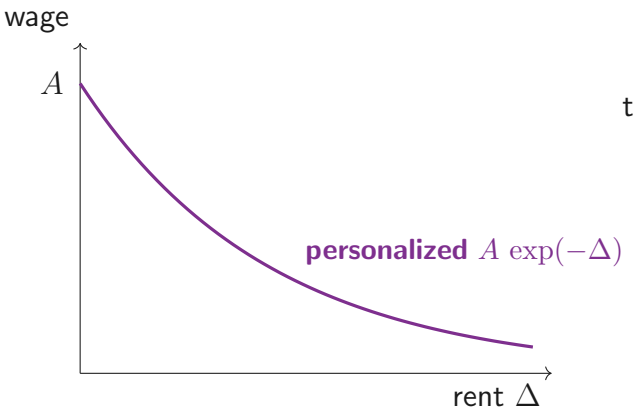
$$\frac{A - w^P}{w^P} = \frac{1}{\lambda_J}, \qquad \lambda_J \equiv \frac{\partial \log P_j}{\partial \log w_j}$$

- λ_J: the elasticity of labor supply to the firm
- exponential tastes, rivals paying w^P:

$$P_j = \frac{x_j}{J}\Big[1 - \max\{0,\, 1 - 1/x_j\}^J\Big], \qquad x_j \equiv \left(\frac{w_j}{w^P}\right)^{1/\sigma}$$

  - wage cut: the bracket is one, so a one percent cut costs $1/\sigma$ percent of hires
  - raise: hires grow by less, because some workers already take its job for sure
  - no kink at $w_j = w^P$: $\lambda_J = 1/\sigma$ for every J, and $w^P = A/(1+\sigma)$

Personalized Wages: Beating the Runner-Up



the runner-up offers $w = A$

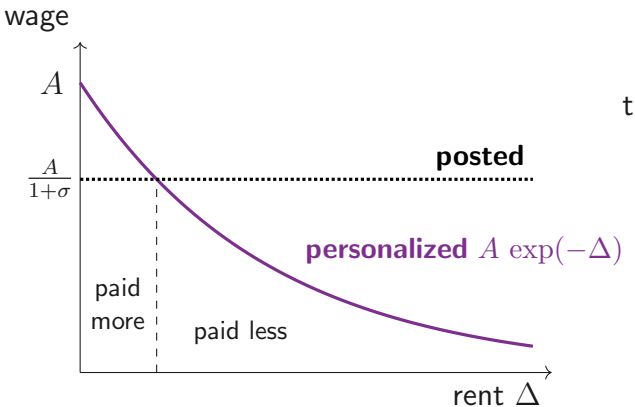
the favorite matches that utility offer

$$\log w^D + \varepsilon_{(1)} = \log A + \varepsilon_{(2)}$$

$$w^D = A \exp(-\Delta)$$

the firm she likes best cuts her wage by exactly her rent

Personalized Wages: Beating the Runner-Up



the runner-up offers $w = A$

the favorite matches that utility offer

$$\log w^D + \varepsilon_{(1)} = \log A + \varepsilon_{(2)}$$

$$w^D = A\,\exp(-\Delta)$$

$$\mathbb{E}\, w^D = \frac{A}{1+\sigma} = w^P$$

a small rent means a wage above the posted wage; the average wage is the same

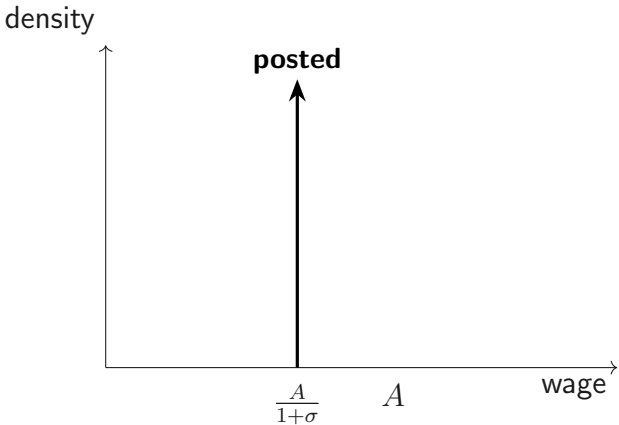
Bidding: A Firm Knows Only Its Own Match

- firm j observes ε_j and offers $w^B(\varepsilon_j)$: a first-price auction for the worker
  - personalized wages are the second-price auction
- a firm with a better match shades its offer by more: w^B is decreasing

$$w^B(\varepsilon) = \mathbb{E}\left[\, w^D \,\middle|\, \varepsilon_{(1)} = \varepsilon \,\right]$$

- the bid is the expected personalized wage, given that the firm's match is the best
  - same average wage
  - the personalized wage adds noise from the runner-up's draw
- risk-neutral firms bid like bidders with constant absolute risk aversion
  - who are indifferent between the two auctions (Matthews 1987)

Same Average Wage, Same Profit

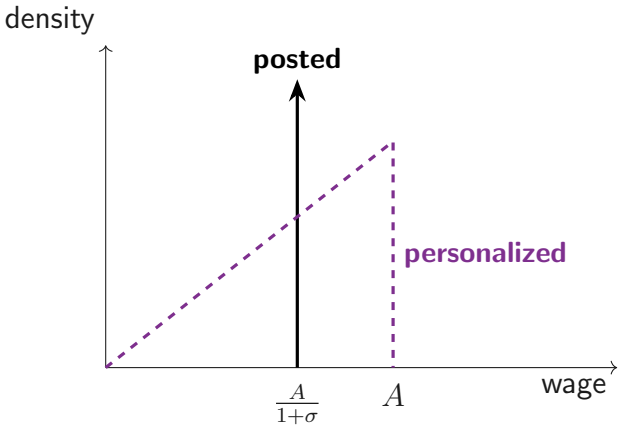


Proposition

With exponential tastes, the expected wage is $A/(1+\sigma)$ in all three regimes, for every $J \geq 2$ and every γ. Expected profit is the same in all three.

posted wages: every worker is paid the same

Same Average Wage, Same Profit

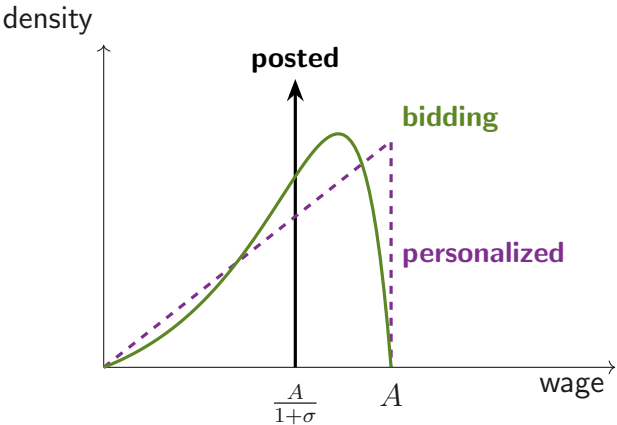


Proposition

With exponential tastes, the expected wage is $A/(1+\sigma)$ in all three regimes, for every $J \geq 2$ and every γ. Expected profit is the same in all three.

personalized wages: anything from nothing to her full product, with the same mean

Same Average Wage, Same Profit



Proposition

With exponential tastes, the expected wage is $A/(1+\sigma)$ in all three regimes, for every $J \geq 2$ and every γ. Expected profit is the same in all three.

bidding: less dispersed than personalized wages, and the same mean again
every regime puts her at her favorite firm, so equal wages mean equal profits

How Special Is the Benchmark?

- general preferences $\phi(u(w) + \varepsilon_j)$ and any log-concave taste density

$\mathbb{E}\, w^D$ relative to w^P, if $1/u'(w)$ is	convex	**linear**	concave
exponential tastes	lower	**the same**	higher
other tastes	either	higher	higher

- linear: $u(w) = \log w$ or w; convex: $u(w) = w^{1-\kappa}/(1-\kappa)$ with $\kappa > 1$
- posted markdown: set against a worker with no rent; personalized: her actual rent
  - log-concave tastes: rents are shorter than the marginal worker suggests
  - log utility, normal tastes, $J = 2$: personalized wages are 16 percent higher

Outline

1. Model

2. Wages and Profits

3. Welfare

4. Employment

5. Conclusion

Equal Wages, Different Risks

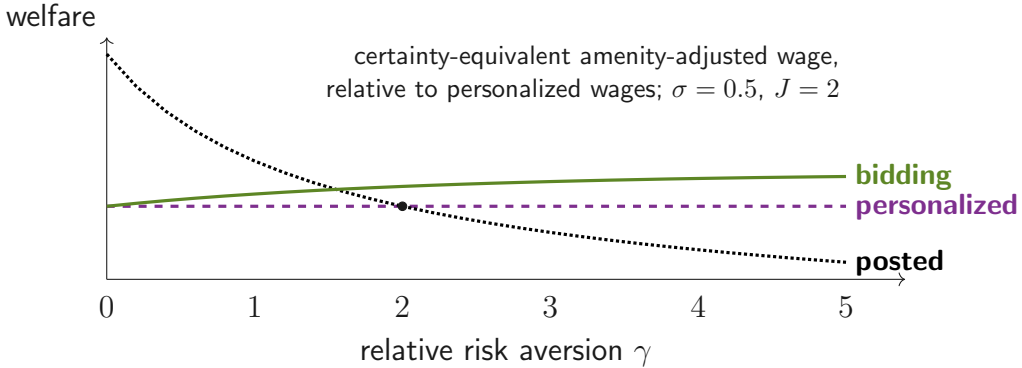
- welfare depends on the wage and the amenity: $U(w, \varepsilon_j) = \dfrac{\left(w \, \exp(\varepsilon_j)\right)^{1-\gamma}}{1-\gamma}$

$$\textbf{personalized}: \quad w^D \exp(\varepsilon_{(1)}) = A \, \exp(\varepsilon_{(2)})$$

$$\textbf{posted}: \quad w^P \exp(\varepsilon_{(1)}) = A \, \exp(\varepsilon_{(2)}) \times \frac{\exp(\Delta)}{1+\sigma}$$

- **personalized**: as if she worked at her second-best firm for her full product
- **posted**: her favorite firm's amenity, less the markdown
  - **an independent random factor** multiplies her amenity-adjusted wage
- the factor's arithmetic and geometric means exceed one; its harmonic mean is one
  - a worker with $\gamma = 2$ values a lottery at its harmonic mean
- the personalized wage is low exactly when her rent is high: it insures her

Personalized Wages Are Never the Worker's First Choice



- she prefers posted to personalized wages if and only if $\gamma < 2$
- she prefers bidding to personalized wages for every $\gamma > 0$
- profits are the same in every regime, so these are Pareto rankings

Outline

1. Model

2. Wages and Profits

3. Welfare

4. **Employment**

5. Conclusion

The Option Not to Work

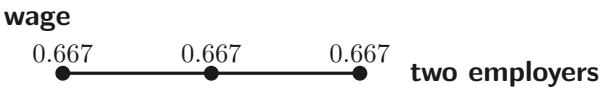
- not working pays $w_0 < A$ with taste ε_0, drawn from the same distribution
  - firms that personalize observe ε_0; firms that post or bid do not

Proposition ($J \geq 2$, exponential tastes, $w_0 \leq w^P$)

The expected wage of an employed worker is $A/(1+\sigma)$ in all three regimes. Unemployment under personalized wages is $(1+\sigma)^{-1/\sigma} < 1$ times unemployment under posted wages.

- the posted markdown leads her to turn down jobs worth more than not working
  - a firm that knows her outside option never loses such a match

Two Employers: More Jobs at the Same Pay

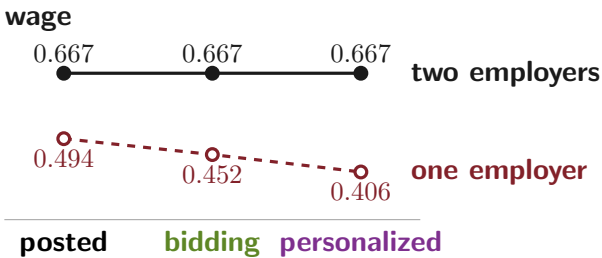




expected wage of an employed worker; $\sigma = 0.5$, $A = 1$, $w_0 = 0.4$

- two employers: information raises employment and leaves pay alone

One Employer: The Textbook Is Right





expected wage of an employed worker; $\sigma = 0.5$, $A = 1$, $w_0 = 0.4$

- two employers: information raises employment and leaves pay alone
- one employer: personalization pushes her down to the value of not working
  - efficient employment, and the firm keeps the surplus
- competition among employers protects her, not the privacy of her preferences

Diminishing Returns: The Labor Share Does Not Move

- each firm produces L^{α}: replace A with the marginal product of labor
- personalized wages raise employment, which lowers the marginal product
  - employed workers earn the same share of a smaller marginal product
- the labor share is $\alpha/(1+\sigma)$ under posted and under personalized wages
  - output, earnings, and profits rise in the same proportion

	employment	wage if employed	output	labor share
posted	0.910	0.578	1	0.444
personalized	0.959	0.568	1.035	0.444

$\alpha = 2/3$, $\sigma = 0.5$, $J = 2$, $w_0 = 0.3$; output relative to posted wages

Outline

1. Model

2. Wages and Profits

3. Welfare

4. Employment

5. Conclusion

Three Regimes, Three Policies

- a ban on personalized pay: **posted wages**
  - the same average wage and lower employment
  - workers gain only if they are not too risk averse: $\gamma < 2$ when everyone works
- a rule against conditioning pay on a worker's other options: **bidding**
  - in the spirit of bans on asking applicants about their pay history
  - when everyone works: better for every risk-averse worker, same profits
- a single employer is the case in which personalization takes the surplus
  - the remedy is a second employer
- what matters: which information firms may use, and whether a rival bids for her

Conclusion

Does income shift toward those who harness the information?

- with competing employers, no: the runner-up disciplines every wage offer

- a paradigm: tastes, market power, information, equilibrium among informed firms
- benchmark: the same average wage and profit whatever firms know
- information moves risk: workers with $\gamma < 2$ prefer posted wages
- information moves employment: more jobs, the same labor share
- prices, in general equilibrium: the same answer, for a different reason

The Distributional and Welfare Implications of Personalized Wages

Robert Shimer
University of Chicago

XXIX Annual Conference of the Central Bank of Chile
November 20, 2026