Discussion of "The Distributional and Welfare Implications of Personalized Wages"

by Robert Shimer

with an eye on the companion paper, "The General-Equilibrium Incidence of Personalized Pricing"

Alejandro Micco
Discussant

XXIX Annual Conference of the Central Bank of Chile

November 20, 2026

What Robert has shown us

A new answer to an old worry

	The textbook (Pigou 1920)	Robert's equilibrium benchmark
What firms learn	Her reservation wage	Her taste for each job
What is held fixed	Every other employer	Nothing: every rival responds
Wages and profits	The employer takes everything	Same average wage $A/(1+\sigma)$, same profits
Risk	—	Personalized pay insures the rent Δ; posted preferred iff $\gamma < 2$, bidding always beats personalized
Employment	Efficient	More jobs, same labor share
When personalization takes the surplus	Always	Only with a single employer (Postel-Vinay and Robin 2002): "the remedy is a second employer"

The runner-up, who will pay her full product, disciplines every offer: **competition among employers protects her, not the privacy of her preferences.**

Shimer (2026a): the conference paper (personalized wages). Shimer (2026b): the companion paper (personalized prices). Talk slides 8, 18, 22, 25.

Why I like this paper

A paradigm, not just a result

A paradigm

Tastes give market power; information is a primitive; every firm is informed and every other firm responds.

Exact

Memorylessness delivers closed forms for every number of employers, not only in the limit.

The auction lens

Employers bid for the worker: posting is price competition, bidding a first-price auction, personalization a second-price auction.

Policy-ready

A ban, a pay-history rule, or a second employer: three policies with clear predictions.

My comments build on this framework: **one connection, one proposal, and two thoughts.**

Shimer (2026a); Matthews (1987); Card et al. (2018); Berger, Herkenhoff and Mongey (2022).

Three thoughts and one proposal

Roadmap

1 **Two papers, one answer**
Personalized wages (this conference) and personalized prices (the companion paper): two mechanisms, a common structure, and the link to Rhodes and Zhou (2024)

2 **A proposal: one vacancy per firm**
"The remedy is a second employer" — when does that employer have room to hire? *Ceteris non paribus* in the labor market

3 **Who gains and who loses**
The companion paper's distributional question, asked of wages; what γ captures

4 **Extensions, both sides of the budget, and Chile**
Timing and information; wages and prices personalized together; Ley 21.431

Two papers, one answer, two mechanisms

	Wages (conference paper)	Prices (companion paper)
Market	J employers, tastes for jobs	Continuum of monopolistic competitors
Why the textbook fails	It holds every other *employer* fixed	It holds every other *price* fixed
What disciplines the firm	The runner-up bids her full product (strategic)	Other firms' prices fall (general equilibrium)
Neutral benchmark	Exponential tastes: same mean wage, same profits	Constant elasticity or Pareto tastes: same allocation
Beyond the benchmark	Log-concave tastes (log or linear utility): personalized wages higher on average	Marshall's second law (e.g. θ with increasing generalized failure rate): poorer workers gain, richer lose
An equivalence	Posted wages plus a mean-zero tax on the rent (next slides)	Uniform prices plus a progressive income tax

Robert's own summary: "the same answer, for a different reason." **The two papers are best read together.**

Shimer (2026a), the conference paper, talk slide 29; Shimer (2026b), companion paper, abstract and Sections 3, 5, 7–8.

The wage paper is the oligopoly channel: Rhodes and Zhou (2024)

With $u(w) = w$, the labor model nests the goods-market model of personalized pricing:

Rhodes and Zhou (2024)	**Shimer (2026a), $u(w) = w$**
Consumer's value v_j	Worker's taste ε_j
Price p	Minus the wage, $-w$
Marginal cost c	Minus marginal product, $-A$
Full market coverage	No outside option
Partial coverage	Nonemployment
$p = c + (v_{(1)} - v_{(2)})$	$w = A - (\varepsilon_{(1)} - \varepsilon_{(2)})$

What the labor version adds

Curvature. A job is most of a household's income; a purchase is not.

A floor. Wages stop at zero or the minimum wage; log utility keeps $Ae^{-\Delta}$ positive.

Risk-averse bidding. Concave u is what makes bidding differ from personalization (Matthews 1987).

Lotteries over jobs. γ separates how jobs are ranked from attitudes to risk.

Rhodes and Zhou (2024), *AER* 114(7): 2141–2170; Shimer (2026a), footnote 12; Shimer (2026b), footnotes 3 and 11.

Neutrality sits at a boundary — in both papers

Robert shows that with log-concave tastes (log utility) personalized wages are **higher** (normal, $J = 2$: $+16\%$).

Rhodes–Zhou find the same in goods markets (Prop. 1; the exponential is their edge case, fn. 21). Companion paper: Pareto tastes are the boundary; Marshall's second law makes poorer buyers gain.

Right: Gumbel tastes with **log** utility, the Card et al. (2018) workhorse, give the same message.

> A suggestion: one common statement for both papers — neutrality at the boundary, gains for the priced side in the regular case — and logit to connect with monopsony, where J matters.

σ	J	Posted wage	Mean personalized	Gap
0.25	2	0.667	0.734	$+10.1\%$
0.25	3	0.727	0.762	$+4.8\%$
0.25	10	0.783	0.791	$+1.0\%$
0.50	2	0.500	0.571	$+14.2\%$
0.50	10	0.643	0.653	$+1.5\%$

Gumbel tastes with scale σ, $u = \log w$, $A = 1$, no outside option. Posted: $(A - w)/w = \sigma J/(J - 1)$. Personalized: $\mathbb{E}[Ae^{-\Delta}]$ with the distribution of Δ in Shimer's Lemma 2.

Discussant's calculation. Rhodes and Zhou (2024), Prop. 1, fn. 21; Shimer (2026a), talk slide 19, Lemma 2; Shimer (2026b), Proposition 1.

An equivalence for wages, in the spirit of the companion paper

COMPANION PAPER (PRICES)

Personalized prices = uniform prices + a progressive income tax

Balanced by a lump-sum tax on, or transfer to, firms. Low-income workers receive a subsidy; high-income workers pay.

CONFERENCE PAPER (WAGES), RESTATED

Personalized wages = posted wages × a mean-zero tax factor

$$w^D = Ae^{-\Delta} = w^P \cdot (1+\sigma)\, e^{-\Delta}$$

$$\mathbb{E}\left[(1+\sigma)e^{-\Delta}\right] = 1$$

Workers with a small rent ($\Delta < \log(1+\sigma)$) are subsidized; workers with a large rent are taxed.

The same language for both papers: **personalization as a tax system** — progressive in income for prices, an insurance-like tax on the rent for wages.

Discussant's derivation from Shimer (2026a), talk slides 16 and 21; Shimer (2026b), Theorem 2 and Proposition 5.

The employment result has a goods-market twin

	Rhodes–Zhou, partial coverage	**Shimer, option not to work**
Mean transfer	Average price paid $= 1 + c$, the uniform price	Expected wage of the employed $= A/(1 + \sigma)$
Participation	Sales rise	Employment rises
Firms	Profits rise	Profits rise: same margin, more hires
The other side	Consumer surplus falls	Linear-ϕ worker ($\gamma = 1$) prefers posted wages
Emphasis	Consumer surplus	Employment

Same mechanics, two emphases. Robert's policy slide already notes that a ban helps workers with $\gamma < 2$; **the abstract could say so as well.**

Exponential tastes. Rhodes–Zhou, Prop. 2 and p. 2155 ($\sigma = 1$); Shimer, talk slides 24–25 and 28.

A natural next step: only the platform personalizes

RHODES AND ZHOU, SECTION III.A

The mixed regime

k firms with data set personalized prices after the other $n - k$ post uniform prices.

At high coverage the mixed regime is the **worst for consumers**: data firms poach rivals' loyal customers, which lowers match quality, and segmentation raises profits.

THE LABOR-MARKET VERSION

Apps and everyone else

Uber, Lyft, DoorDash and Instacart see acceptance histories; restaurants, shops and warehouses post wages.

Robert's framework is ideal to ask: does the runner-up still discipline the app when the runner-up posts?

A natural extension: **a share α of informed employers**. In the companion paper (and in Rhodes–Zhou) the mixed case is the worst: welfare is single-troughed in α.

Rhodes and Zhou (2024), Section III.A; Shimer (2026b), Sections 3.5 and 6.3; FTC orders of October 2026 (Shimer 2026a, talk slide 2).

A proposal: one vacancy per firm

Robert's policy slide

"A single employer is the case in which personalization takes the surplus — **the remedy is a second employer**."

My proposal, in the spirit of the companion paper

Ceteris non paribus in the labor market: when one firm fills its vacancy, it stops being the runner-up for the next worker. Ask when the second employer has **room to hire**.

The paper follows the discrete-choice tradition in labor (Card et al. 2018; Berger et al. 2022): tastes for firms and constant returns, so a firm can always take one more worker. In labor markets the vacancy is the unit of hiring, and that feature matters even more than in goods markets.

Shimer (2026a), talk slide 28; Shimer (2026b), introduction; Card et al. (2018); Berger, Herkenhoff and Mongey (2022).

Why one vacancy matters even more in labor than in goods

Discrete choice, à la Card et al.

Workers choose among firms; each firm faces a smooth labor supply; constant returns separate its problem worker by worker.

Both sides are indivisible

A worker takes one job and a position takes one worker: a one-to-one matching market. A seller of goods usually sells many units.

Hiring is for positions

A filled firm's offer to another worker is A minus the profit on the worker it would let go — not A.

Three questions organize the proposal

Who competes with whom?

Does capacity bind?

Who makes the offers?

Shimer (2026a), footnote 2 and Section 7.

Who competes with whom: three polar cases

(A) ONE WORKER, J FIRMS

Firms compete

This is Robert's model. The runner-up offers A, so $u(w) + \varepsilon_{(1)} = u(A) + \varepsilon_{(2)}$.

The worker gets her second-best job at full pay.

(B) ONE JOB, N WORKERS

Workers compete

The firm offers each candidate her reservation wage and hires the best fit.

She gets exactly her outside option, for any N. Product twin: one seat, N buyers — full extraction.

(C) N WORKERS, J ONE-JOB FIRMS

An assignment game

A worker's floor is the best offer from an idle vacancy, if one exists; otherwise her outside option.

What matters is the number of idle vacancies, not the number of firms.

What protects the worker: **a second employer with room to hire.**

Case (b) generalizes Shimer's single-employer result in Section 6.

One vacancy, N candidates: personalization fills the job and takes all the rent

	Posted wage			Personalized offers		
N	Wage	Filled	Rent	Wage	Filled	Rent
1	0.494	0.672	0.375	0.406	0.920	0
2	0.443	0.834	0.215	0.314	0.994	0
3	0.412	0.895	0.150	0.263	0.999	0
5	0.358	0.923	0.093	0.209	1.000	0
10	0.283	0.944	0.047	0.153	1.000	0

Shimer's benchmark: log utility, exponential tastes with $\sigma = 0.5$, $A = 1$, outside option $\log 0.4 + \varepsilon_0$. Rent: utility above the outside option, per candidate. $N = 1$ reproduces Shimer's single-employer numbers.

Efficiency gain

The best-fitting candidate is always hired, and the job is almost always filled.

Zero rent, every state

So every ϕ prefers posted wages. Risk attitudes stop mattering.

Discussant's calculations in Shimer's Section 6 specification.

Many workers, many one-job firms: an assignment game

- Full information; firms raise an offer only when outbid. Firm-proposing salary auctions (Crawford–Knoer 1981; Kelso–Crawford 1982; Demange–Gale–Sotomayor 1986).
- Outcome: the **firm-optimal stable assignment** — the lowest competitive wages.
- Shimer's Proposition 2 is its one-worker case.
- The core is an interval: Böhm-Bawerk's horse market (Shapley–Shubik 1971). Even with one worker, Shimer's personalized wage is its **bottom**.

Runner-up firm's best offer $= A -$ profit on the worker it would otherwise hire

Worker's floor $= \max\{$best offer from an idle vacancy, outside option$\}$

The relevant count is **idle vacancies per worker**, not firms per worker.

Shapley–Shubik (1971); Crawford–Knoer (1981); Kelso–Crawford (1982); Demange–Gale–Sotomayor (1986).

Small markets: with as many jobs as workers, wages fall to the monopsony level

Mean wage of employed workers under personalized offers

Workers	1 job	2 jobs	3 jobs	4 jobs
1	0.406	**0.667**	0.667	0.667
2	0.313	**0.413**	**0.635**	0.646
3	0.262	0.309	**0.413**	**0.614**
Shimer	0.406	0.667	0.667	0.667

Rows: number of workers N; columns: number of one-job firms J. Shimer: room to hire. Orange: as many jobs as workers. Blue: one idle vacancy. Row $N = 1$ is Shimer's Props. 10–11; other rows simulated (6,000 draws per cell; standard errors about 0.002). All cells: the **firm-optimal** stable outcome, the selection behind Shimer's personalized wage; other stable wages are higher, at most A.

0.41

With as many jobs as workers the mean wage is about the single-employer level (0.406), not Shimer's 0.667, even though every worker has several firms.

One idle vacancy restores most of it: 0.64 with two workers, 0.61 with three.

Discussant's simulations; $\sigma = 0.5$, $A = 1$, $w_0 = 0.4$; firm-proposing ascending auction.

Large market: wages jump when vacancies outnumber workers

Mean wage of employed workers, personalized offers



0.19 → 0.51

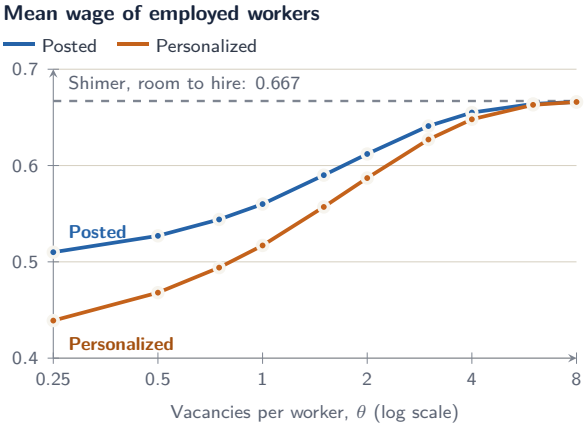
Mean wage as vacancies per worker go from 0.9 to 1.1. The short side takes the surplus: the "stark effect of competition" (Ashlagi, Kanoria and Leshno 2017).

With many idle vacancies the wage approaches Shimer's 0.667. With fewer jobs than workers it falls far below his single-employer level.

Discussant's simulations: $N = 40$ workers, 20–40 draws per point, firm-proposing ascending auction; $\sigma = 0.5$, $A = 1$, $w_0 = 0.4$.

With frictions: each firm can make one offer

2 · A proposal: one vacancy per firm



Each one-job firm contacts one worker at random; a worker hears from $k \sim \text{Poisson}(\theta)$ firms.

$k \geq 2$: Bertrand between firms, Shimer's neutrality applies.

$k = 1$: personalization pushes her to her outside option.

At $\theta = 1$: wage 0.560 posted vs 0.517 personalized (-8%); employment 0.50 vs 0.59.

Julien, Kennes and King (2000); Albrecht, Gautier and Vroman (2006). Discussant's simulations.

Welfare: the $\gamma = 2$ threshold with one-vacancy firms and frictions

Vacancies per worker, θ	Hires with a single offer	Worker prefers posted wages when
0.5	77%	every γ from 0 to 12
1	57%	every γ from 0 to 12
2	30%	every γ from 0 to 12
4	7%	every γ from 0 to 12
8	0.3%	$\gamma < 2.15$, or $\gamma > 5.85$
Room to hire (Shimer)	0%	$\gamma < 2$

Urn-ball model of the previous slide; CRRA γ over the amenity-adjusted wage, as in Shimer's Prop. 7. Share of personalized hires whose employer had no rival.

Insurance needs a runner-up. With one offer, the personalized wage pushes the worker to her outside option instead of insuring her.

With few rival offers ($\theta \leq 4$), **every worker** — however risk-averse, up to $\gamma = 12$ — prefers posted wages.

Discussant's simulations; $\sigma = 0.5$, $A = 1$, $w_0 = 0.4$.

For a central bank: personalized wages move more with tightness

POSTED WAGE

+16%

0.527 → 0.612 as θ goes from 0.5 to 2

PERSONALIZED WAGE

+25%

0.468 → 0.587 over the same range

If this carries over, algorithmic pay implies

- A steeper wage Phillips curve: wages respond more to slack
- A more procyclical labor share when capacity binds in downturns
- Harm concentrated in recessions, exactly when the runner-up employer is missing

Urn-ball model; $\sigma = 0.5$, $A = 1$, $w_0 = 0.4$. A conjecture to test, not a result.

Building on Section 7: can a very concave $Y(L)$ mimic a firm with one vacancy?

$$Y_k(L) = \frac{A}{k}\left[\log\left(1+e^{k}\right) - \log\left(1+e^{k(1-L)}\right)\right]$$

$$Y_k'(L) = \frac{A}{1+e^{k(L-1)}}, \qquad Y_k''(1) = -\frac{Ak}{4} \to -\infty, \qquad Y_k(L) \to A\min\{L,1\}$$

CAPACITY SLACK

Shimer survives

Labor per firm below L^*: $Y' \to A$, and the constant-returns results hold.

CAPACITY BINDS

Determinate, and costly

With an outside option, employment is capped and wages adjust instead (two slides ahead).

EXACTLY AT CAPACITY

Indeterminate

With no outside option, the limit wage depends on which smooth family you pick (next slide).

Normalization $L^* = 1$. Curvature at one point is not enough: the whole family matters. Shimer (2026a), Section 7.

At the kink, the limit wage is an artifact of the approximation

No outside option, labor per firm exactly L^*. Shimer's formula gives mean wage $Y'(L^*)/(1+\sigma)$.

Family above: $Y'_k(L^*) = A/2$ for every k
$\Rightarrow$ wage $\dfrac{A}{2(1+\sigma)}$

Alternative: $Y'_k(L) = \dfrac{A}{1 + c\, e^{k(L-L^*)}}$
same limit $A\min\{L, L^*\}$, but $Y'_k(L^*) = A/(1+c)$
$\Rightarrow$ any wage in $\big(0,\ A/(1+\sigma)\big)$

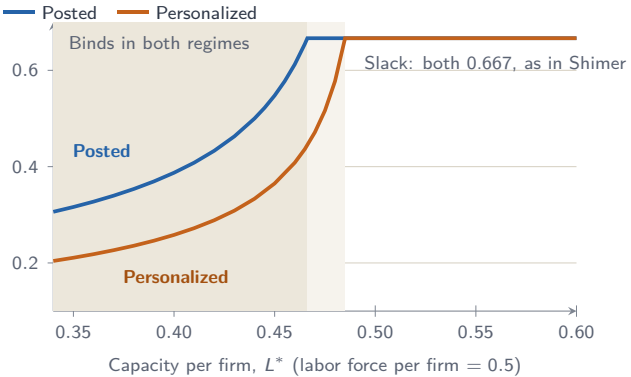
At the kink the marginal product is any number in $[0, A]$. A smooth family picks one — just as a matching protocol picks one point of the core.

Production limits do not determine wages.
The allocation rule does.

Böhm-Bawerk's horse market: with equal numbers on both sides, the price is indeterminate within the core.

With an outside option, binding capacity turns the employment gain into a wage cut

Mean wage of employed workers, $k \to \infty$



When capacity binds, employment $= JL^*$ in both regimes:

posted:
$w_0 \left[(J+1)(1-JL^*)\right]^{-\sigma}$

personalized:
posted $/\,(1+\sigma)$

The labor share falls by the same factor. At $L^* = 0.45$: $0.548 \to 0.365$. Shimer's Cobb–Douglas example: $0.578 \to 0.568$.

Discussant's calculation with Shimer's eq. (20); $J = 2$, $\sigma = 0.5$, $A = 1$, $w_0 = 0.3$ (his Section 7 example).

How the paper's results extend to one-vacancy firms

Result	Room to hire (the paper)	One job per firm, few idle vacancies per worker
Mean wage of employed	Same in all regimes	Lower under personalization: -8% at $\theta = 1$ with frictions; -33% when capacity binds
Employment	Higher under personalization	Still higher with frictions; no gain when capacity binds
Welfare threshold	Posted preferred iff $\gamma < 2$	Posted preferred for every $\gamma \leq 12$ when $\theta \leq 4$
Labor share	Unchanged with Cobb–Douglas	Falls by $1/(1+\sigma)$ when capacity binds
Efficiency	All regimes assign the worker to her favorite firm	Personalization fills more jobs and hires the best fit
Who is protected	Workers with a second employer	Workers in tight markets

Capacity rows: fewer jobs than willing workers. Friction rows: urn-ball, θ vacancies per worker; many workers get a single offer. Numbers from the preceding slides.

How the proposal could fit in the paper

1 · A one-vacancy section

The assignment-game version of "beating the runner-up": personalized wages as the firm-optimal stable wages.

2 · Tightness as a key statistic

Wages, employment and welfare by vacancies per worker, alongside the number of employers J.

3 · A posted benchmark exists

Robert's own model of coordination frictions (Shimer 2005) — one vacancy per firm, posted wages, workers who cannot coordinate applications — is the natural posted side; adding taste shocks would complete it.

4 · A refined policy message

"The remedy is a second employer" — **with room to hire**. Tight labor markets protect workers.

Shimer (2005), "The Assignment of Workers to Jobs in an Economy with Coordination Frictions," *JPE* 113(5).

Who gains? The companion paper asks it; the wage paper could too

Companion paper: personalization = segmentation by group (third-degree) + personalization within group. Segmentation lowers prices for elastic buyers (the poor) and raises them for inelastic ones (the rich).

The labor analogue: groups differ in attachment σ. Posted wages cannot condition on the group; personalized offers can.

> In goods, the inelastic side is the rich: progressive. In labor, the inelastic suppliers are the attached — less mobile, caregivers, rural: **possibly regressive**.

Group	σ	Pooled posted	Personalized	Change
Footloose	0.25	0.727	0.800	+10%
Attached	0.75	0.727	0.571	−21%
Average		0.727	0.686	−5.7%

Pooled posted wage: $(A-w)/w = 1/\bar{\lambda}$ with $\bar{\lambda}$ the average of $1/\sigma_g$. Within each group the personalized mean is $A/(1+\sigma_g)$, as in Robert's benchmark, so the whole effect is segmentation. Profits rise: with heterogeneous σ, profit neutrality fails.

Discussant's calculation: exponential tastes within groups, $J \geq 2$, $A = 1$, no outside option, equal group sizes. Shimer (2026b), Propositions 3–4 and Theorem 1.

From ex ante to interim, and what γ captures

3 · Who gains and who loses

DISTRIBUTION

From ex ante to interim

Robert's welfare comparison is ex ante, before tastes are drawn. Interim, personalization moves pay from workers attached to one employer (large Δ) to footloose ones.

If attachment differs by group — women trade wages for shorter commutes (Le Barbanchon, Rathelot and Roulet 2021) — equally productive workers get different pay.

Suggestion: report outcomes by Δ and by group (previous slide).

INTERPRETATION

One more reading of γ

Robert notes that γ governs lotteries over jobs, not which job she takes. With $U = \frac{(we^{\varepsilon})^{1-\gamma}}{1-\gamma}$, the cross-partial $\partial^2 U / \partial w \, \partial \varepsilon$ has the sign of $1-\gamma$: for $\gamma > 1$, money is worth less in a job she likes.

Evidence on state-dependent utility points to complementarity (Finkelstein, Luttmer and Notowidigdo 2013), and saving across job spells lowers the curvature per spell. Both suggest $\gamma < 2$, where posted wages are preferred.

Shimer (2026a), talk slides 11 and 21–22.

Timing, information and productivity: three extensions

TIMING

Simultaneous offers are the best case

Robert contrasts his model with sequential auctions (Postel-Vinay and Robin 2002). With sequential or per-trip offers, a hire from unemployment often faces one employer at a time.

INFORMATION

Data come from workers

Acceptance histories invite strategic rejection and ratchet effects. Disclosure choice (Ali, Lewis and Vasserman 2023) and information design (Bergemann, Brooks and Morris 2015) fit Robert's "which information firms may use."

PRODUCTIVITY

Heterogeneous firms

With different A across firms, posted wages send some workers to the wrong firm and personalization corrects it: a surplus gain, and a question of who gets it.

Shimer (2026a), talk slides 5 and 28.

Both sides of the household budget

EACH PAPER HOLDS ONE SIDE FIXED

Wages or prices, not both

The wage paper has no goods market; the price paper has a competitive labor market with one wage.

PLATFORMS DO BOTH

Drivers and riders

A ride-hailing app personalizes what it pays the driver and what it charges the rider.

FOR A CENTRAL BANK

Household-specific real wages

With personalized prices the price index differs by household; real wages need both sides.

A joint question: **does the progressivity on the price side offset the extraction or insurance on the wage side?**

Shimer (2026a): no goods market; Shimer (2026b): a competitive labor market with a common wage.

Chile: a natural laboratory for the paper's predictions

Ley 21.431

Regulates platform work, in force since September 2022, with algorithmic-transparency duties.

About 3% of employment

Most of it, 83%, in passenger transport; drivers multi-home across Uber, DiDi, Cabify and InDriver.

Three testable predictions

1. **Multi-homing premium.** A driver active on two apps has a runner-up: her personalized offers should be better than a single-app driver's.
2. **Tightness gradient.** The gap between personalized and posted pay should be larger where vacancies per worker are low — across cities and over the cycle.
3. **Outside-option pass-through.** With competing employers, pay should not track unemployment benefits (cf. Jäger et al. 2020); with one employer it should, one for one.

Sources: Diario Constitucional (2026); Centro Cultural de la Cooperación; RIAL (2025).

Two interesting papers, and a proposal

Summary

1 **Two papers, one answer.** Wages (strategic runner-up) and prices (general equilibrium) share a structure: neutrality at a boundary, a sign beyond it, an equivalence as a tax system. Cite Rhodes–Zhou in both.

2 **A proposal: one vacancy per firm.** *Ceteris non paribus* in the labor market: "the remedy is a second employer" — with room to hire. When idle vacancies are scarce personalization lowers wages (-8% to -33%); tightness becomes the key statistic.

3 **Who gains and who loses.** The companion's incidence question, asked of wages: segmentation by attachment may make personalized pay regressive; γ also as complementarity.

4 **Extensions and Chile.** Timing, endogenous data, both sides of the budget; Chile's platform market as a laboratory.

Thank you, Robert, for a wonderful paper

Backup slides follow: closed form, methods and data.

The capacity limit in closed form

Backup

Shimer's eq. (20), log utility, exponential tastes, $J \geq 2$:

$$e = 1 - \frac{1}{J+1}\left(\frac{\kappa\, w_0}{Y'(e/J)}\right)^{1/\sigma}, \quad \kappa = 1+\sigma \text{ posted}, \ \kappa = 1 \text{ personalized}$$

As $k \to \infty$, $Y'_k \to A$ below L^* and $\to 0$ above it. If JL^* is below the constant-returns employment rate, capacity binds:

$$e = JL^*, \quad Y' = \kappa\, w_0 \left[(J+1)(1-JL^*)\right]^{-\sigma},$$
$$\mathbb{E}[w \mid \text{employed}] = Y'/(1+\sigma)$$

So personalized $\div$ posted $= 1/(1+\sigma)$. Prop. 11 needs $w_0 \leq w^P$, i.e. $L^* \geq 1/(J+1)$.

Check with the smooth family, $L^* = 0.45$

	Posted		Personalized	
k	e	wage	e	wage
5	0.791	0.379	0.885	0.340
100	0.879	0.497	0.897	0.360
10,000	0.900	0.547	0.900	0.365
Limit	0.900	0.548	0.900	0.365

$J = 2$, $\sigma = 0.5$, $A = 1$, $w_0 = 0.3$. Thresholds: posted binds below $L^* = 0.466$, personalized below 0.485.

Discussant's derivation from Shimer (2026a), Propositions 11–12.

How the one-job numbers were computed

Backup

Benchmark

$u = \log w$; tastes Exponential, $\sigma = 0.5$; $A = 1$; outside option $\log w_0 + \varepsilon_0$; $w_0 = 0.4$ (Section 6) or 0.3 (Section 7).

Validation

Reproduces Shimer's $J = 1$ numbers (posted 0.494 and 0.672; personalized 0.406 and 0.920) and $J = 2$ (employment 0.880 and 0.946).

Assignment game

Firm-proposing ascending auction from reservation wages; an outbid firm raises its offer 0.2% above the best held offer. $N = 40$.

Frictions

Number of firms reaching a worker $\sim$ Poisson(θ). Posted wage: symmetric best response solved numerically by quadrature; 0.8 million simulated workers per point.

Code available on request.

Data behind the one-job charts

Backup

Assignment game, personalized

θ	Mean wage	Employed
0.25	0.086	0.250
0.5	0.097	0.500
0.9	0.192	0.900
1.0	0.365	0.988
1.1	0.511	0.991
1.5	0.600	0.998
2	0.620	0.999
3	0.645	1.000
5	0.649	0.999

One offer per firm (urn-ball)

	Mean wage		Employed	
θ	Posted	Personalized	Posted	Personalized
0.25	0.510	0.439	0.156	0.205
0.5	0.527	0.468	0.290	0.365
1	0.560	0.517	0.497	0.590
2	0.612	0.587	0.738	0.817
4	0.655	0.648	0.897	0.945
8	0.666	0.666	0.955	0.980

θ = vacancies per worker. Mean wage of employed workers. $\sigma = 0.5$, $A = 1$, $w_0 = 0.4$.

Discussant's simulations in Shimer's Section 6 specification.