


The Churn Ladder

Lawrence Warren David Wiczer

U.S. Census Bureau / CES ▪ Federal Reserve Bank of New York & IZA

April 9, 2026

Workers Climb a Wage Ladder and Descend a Churn Ladder



- Within a **continuous employment spell**, track firm characteristics across job-to-job moves
- **Log churn falls ≈ 40 log points** over 7 transitions; log earnings rise ≈ 30 log points
- Firm **size** is nearly flat—the ladder is about wages and stability, not scale

This Paper

Question: Why is the churn ladder so steep, and what does it imply for labor market power?

1. Document a “churn ladder”: firm-level churn declines steeply with wages, both cross-sectionally and within workers’ job sequences
2. The canonical job-ladder model cannot generate this steepness— **not poaching alone**
3. Separation heterogeneity with endogenous sorting matches churn & tenure
4. Ignoring nonemployment separations **overstates monopsony**/ understates labor elasticity

Data: LEHD Matched Employer-Employee Panel

Source

Census Bureau Longitudinal Employer-Household Dynamics (LEHD)
17-state sample, 1998–2013
(42% of national employment)

Sample

15% random sample of workers
Ages 25–65; firms ≥ 10 employees
1.2 million job-to-job transitions

Key measures

- **Firm churn** = excess hires above net growth ($\equiv \min(H, S)$)
- **Firm wages** = log avg. real quarterly earnings
- **Job-to-job transition** = dominant employer changes with overlapping earnings
- **Employment spell** = continuous employment across jobs (“employment cycle” per Wolpin 1992)

Measuring the Worker Churn Ladder

Strategy: Follow workers within a continuous employment spell; compare firm characteristics at each successive job

Regression on job-transition observations:


$$X_{ij} = \beta_0 + \sum_{k=1}^{N_g} \beta_k \mathbf{1}_{j=\text{job}_k} + \gamma_g + \varepsilon$$

where γ_g = employment-spell fixed effect, and $X_{ij} \in \{\bar{w}_{ij}, \bar{L}_{ij}, \text{EAR}_{ij}\}$ (avg. earnings, firm size, churn rate)

Key features:

- Spell FE absorbs cross-worker heterogeneity and starting-point randomness
- Coefficients β_k = change relative to first job in spell
- 1.2 million job-to-job transitions; spells require ≥ 6 quarters of tenure

Workers Climb Wages, Descend Churn



- Churn falls ≈ 40 log points across 7 transitions; earnings rise ≈ 30 log points
- Pattern holds across sex, race, age, education, firm size, industry (Appendix B)

Why Is the Churn Ladder So Steep?

The canonical explanation: Higher wages \Rightarrow fewer workers are poached \Rightarrow lower churn.

But two very different data-generating processes could produce this:

Firm-side only

Firms jointly set wages and retention policies.
Conditional on wage, all workers face the same separation risk.

\Rightarrow Tenure $\sim \text{Exp}(\delta(w))$ within each wage bin

\Rightarrow No excess dispersion in tenure

Worker heterogeneity

Workers differ in latent δ_i ; low- δ types survive longer spells and sort to high-wage firms.

\Rightarrow Excess mean and variance of tenure within wage bins

\Rightarrow Endogenous sorting amplifies the gradient

This distinction matters

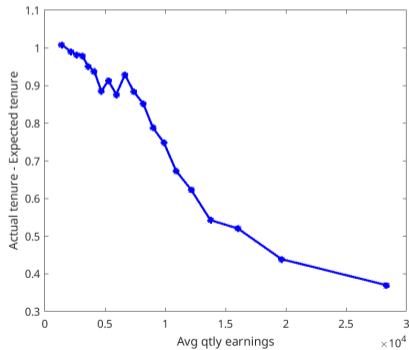
If worker heterogeneity drives the gradient, then firm-focused measures of separation elasticity are mismeasured—with direct implications for monopsony.

Formal Tests: Reject Homogeneous Workers Within Wage Level

First-moment test

Heterogeneity $\Rightarrow E[\tau|w] > \frac{1}{\hat{\delta}(w)}$

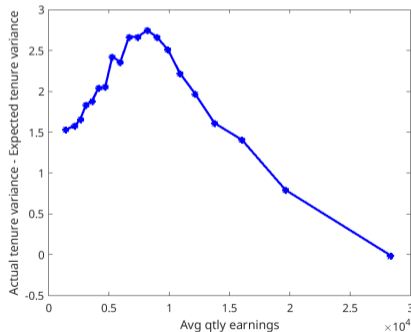
Test stat: $2\hat{\delta}_i(\hat{\tau}_i - \frac{1}{\hat{\delta}_i}) \sim \chi^2$



Second-moment test

Heterogeneity $\Rightarrow \text{Var}(\tau|w) > \frac{1}{\hat{\delta}(w)^2}$

Test stat bootstrapped



Both tests reject uniformity. Suggest more heterogeneity at the bottom (endog selection)

Job Ladder Framework: Burdett-Mortensen Baseline

Churn rate at wage w (hires = separations in steady state):

$$c(w) = \delta + \lambda_e(1 - F(w))$$

- $c'(w) = -\lambda_e f(w) \leq 0$ —churn always falls with wages
- Source: poaching. High-wage firms better retain workers from outside offers
- But δ is a **constant floor**—churn is bounded below by the nonemployment separation rate

The problem

Constant δ is a floor on churn even at the highest wages: Poaching alone generates far too little dispersion.

Two Mechanisms: Firm Policy + Worker Heterogeneity

We allow δ to vary across **both workers and firms**:

1. Firm-level separation policy

Firms bundle wages with job stability. Separation rate scales with wage:

$$\delta_{\text{firm}}(w) = e^{\varepsilon_{\delta} w}$$

$\varepsilon_{\delta} < 0 \Rightarrow$ high-wage firms directly retain workers better

2. Worker heterogeneity + sorting

Workers draw a permanent type $\delta_i \in \{\delta_1, \delta_2\}$. Low- δ workers survive longer spells, make more J2J transitions, and **sort toward high-wage firms endogenously**.


\Rightarrow Worker composition itself steepens the churn gradient

Full churn expression:

$$c(w) = e^{\varepsilon_{\delta} w} \int_0^{\bar{\delta}} s dD(s|w) + \lambda_r + \lambda_e(1 - F(w))$$

λ_r = “reallocation shock” (Godfather offer) to match observed transitions to dominated jobs

Endogenous Sorting Amplifies the Churn Gradient



- At low-wage firms, the workforce is mostly high- δ workers who rarely stay employed long enough to climb the ladder


Calibration

Take empirical $G(w)$, solve for $F(w)$, joint type distribution $h(w_i, \delta_j)$, and minimize distance from LEHD:

1. Worker churn ladder (slope across job transitions)
2. Share of workers with tenure < 1 year across wage deciles
3. Fraction of separations to nonemployment vs. J2J across wage deciles
4. Churn rate across the wage distribution

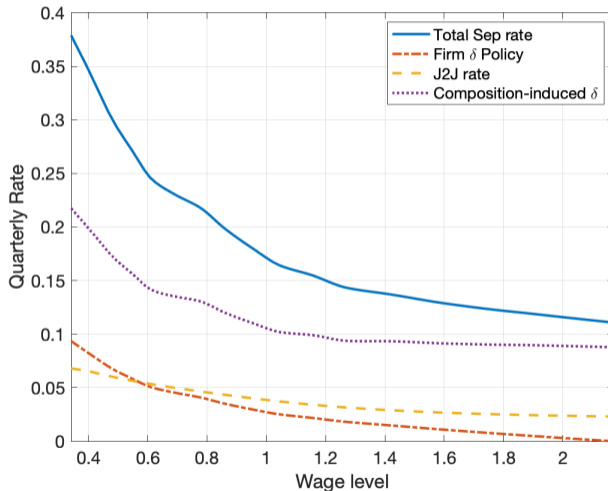
Parameter	Value	Target / Source
λ_u	0.084	Quarterly NE rate
λ_e	0.045	Worker churn ladder
λ_r	0.023	Worker churn ladder
$\varepsilon_{\delta w}$	-0.195	Worker churn ladder
δ_1 (low type)	0.062	Tenure distribution
δ_2 (high type)	0.730	Tenure distribution
p_1 (share of low- δ type)	0.80	Ahn et al. (2023)

Baseline Model Fits the Joint Pattern



- Model replicates the full firm churn gradient—high churn at low-wage firms, near-zero at the top

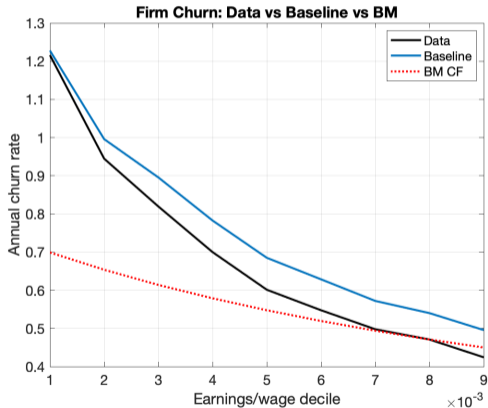
Decomposing the Churn Gradient



- Total churn declines ≈ 25 pp from bottom to top of the wage distribution

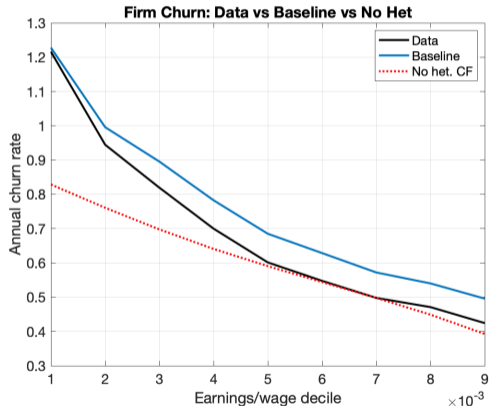
Too Little Heterogeneity in the Canonical Model

Burdett-Mortensen (constant δ)



Constant δ is a floor—churn can't fall enough.

No worker heterogeneity (only $\delta_{firm}(w)$)



Without hetero- δ , tenure and churn inconsistent

Labor Market Power with Constant δ

Standard approach (Manning 2003; Bassier, Dube & Naidu 2021):

Separation elasticity stems from poaching \Rightarrow labor supply curve & monopsonistic markdown

Omits $\delta(w)$ elasticity

- Job-to-job flows (poaching) and treat separations to **nonemployment as wage-invariant** understates the elasticity
- But $\delta'(w) < 0$
- By choosing a higher wage, less poaching & **less quitting**

Full Labor Supply Elasticity with $\delta'(w) \neq 0$

Allowing $\delta'(w) < 0$, the labor supply elasticity becomes:

$$\varepsilon_L(w) = w \left[\frac{\lambda_e(1-u)g(w)}{\lambda_u u + (\lambda_r + \lambda_e G(w))(1-u)} + \frac{\lambda_e f(w) - \delta'(w)}{\delta(w) + \lambda_r + \lambda_e(1-F(w))} \right] \quad (5.3)$$

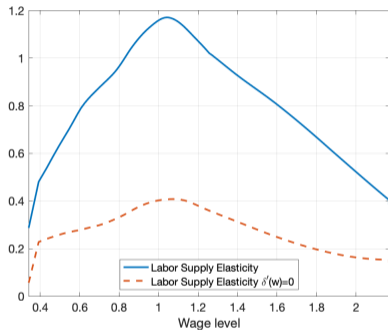
Because $\delta'(w) < 0$, the term $-\delta'(w) > 0$ —the second term is **strictly larger** than $\delta'(w) = 0$.

Special case: suppose $\delta(w) = \kappa(\lambda_r + \lambda_e(1-F(w)))$, $\delta'(w) = -\kappa\lambda_e f(w)$

$$\begin{aligned} \varepsilon_L(w) &= w \left[\frac{\lambda_e(1-u)g(w)}{\lambda_u u + (\lambda_r + \lambda_e G(w))(1-u)} + \frac{\lambda_e f(w)}{\lambda_r + \lambda_e(1-F(w))} \right] \\ &\neq w \left[\frac{\lambda_e(1-u)g(w)}{\lambda_u u + (\lambda_r + \lambda_e G(w))(1-u)} + \frac{\lambda_e f(w)}{\delta + \lambda_r + \lambda_e(1-F(w))} \right] \end{aligned}$$

Bias $f(w) \times \lambda_e \delta A(w)$ where $A' > 0$ **increases with w if $f' \geq 0$** (BM case)

The Elasticity Gap: 2–3× Larger Once We Account for $\delta'(w)$



- **Solid blue:** baseline elasticity with $\delta'(w) < 0$, **Dashed red:** elasticity holding δ fixed
- Gap is largest in the **middle of** w , where the *slopes* of offers density and $\delta(w)$ are *steepest*
- Omitting this channel means the implied markdown is **too large throughout the wage distribution**

Implications for Measuring Monopsony Power

What existing studies do

- Estimate the elasticity of job-to-job separations w.r.t. wages (event studies, quasi-experiments)
- Treat EN separations as exogenous / wage-invariant
- Infer monopsony markdown from this partial elasticity

What they get wrong

- **Understate** total separation elasticity
- EN separations also fall with wages—omitting them **understates how responsive retention is to wages**
- \Rightarrow **Overstate** implied markdown / monopsony power

Bottom line

EE-only estimates of the separation elasticity are **lower bounds**. Incorporating wage-sensitive exits to nonemployment moves implied markdowns toward the more competitive benchmarks found in reduced-form studies.

Conclusion

1. **The churn ladder is a robust empirical fact.** Workers move from high-churn, low-wage firms to low-churn, high-wage firms with each job transition— ≈ 40 log points over 7 moves.
2. **Poaching alone cannot explain it.** Separations to nonemployment are also strongly negatively correlated with wages; constant- δ models miss this entirely.
3. **Worker heterogeneity is essential.** Two formal tests reject homogeneous separation rates within wage bins. Worker heterogeneity accounts for \approx half the churn gradient via endogenous sorting; firm policy explains most of the rest.
4. **Monopsony is overstated.** Once wage-sensitive exits to nonemployment are included, the labor supply elasticity is $2\text{--}3\times$ larger—implying substantially lower markdowns than EE-only approaches suggest.