The Predictability Premium: Fair Workweek Laws, Worker Retention, and the Hiring Freeze

APEP Autonomous Research* @SocialCatalystLab

March 22, 2026

Abstract

When cities mandate predictable schedules for service-sector workers, do firms stabilize their workforce or freeze out new entrants? Using a triple-difference design that compares covered industries (food services and retail) to uncovered industries (manufacturing and professional services) in 45 treated counties versus 3,136 controls, I find that predictive scheduling laws reduce the separation rate by 0.86 percentage points and the new hire rate by 1.09 percentage points in covered industries. Employment rises by 3.1%, consistent with a "retention dividend" that more than offsets reduced hiring inflows. A placebo test on uncovered industries yields a precise zero ($p = 0.86$). These results suggest scheduling mandates create a two-tier labor market: incumbent workers gain stability, while prospective entrants face a thinner hiring pipeline.

JEL Codes: J23, J32, J63, J88, K31

Keywords: predictive scheduling, fair workweek, worker turnover, hiring, labor regulation

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 21m).

1. Introduction

A restaurant worker in New York City learns her schedule for next week at 11 p.m. on Saturday night. By Sunday morning, she has rearranged childcare, canceled a doctor's appointment, and texted a friend to cover a shift she cannot make. This is the daily reality for millions of American service-sector workers whose schedules are set with little advance notice and changed without warning. Beginning with San Francisco in 2015, a wave of "fair workweek" ordinances has sought to end this instability by requiring employers to post schedules 14 days in advance and pay premiums for last-minute changes.

These laws intervene on a margin the labor economics literature has largely overlooked. While minimum wages, overtime rules, and paid leave mandates have been studied extensively (Dube, 2019; Neumark and Wascher, 2008; Meer and West, 2016; Gruber, 1997; Rossin-Slater et al., 2013), the *scheduling* dimension of low-wage work remains empirically underexplored despite its centrality to worker welfare. Schneider and Harknett (2019) document that schedule instability causes income volatility, sleep disruption, and childcare breakdowns, but the causal effects of mandating predictability on labor market flows—hires, separations, and employment—have not been estimated at scale.

This paper provides the first administrative-data estimates of how predictive scheduling laws reshape service-sector labor markets. I exploit the staggered adoption of fair workweek ordinances across six U.S. cities and one state between 2015 and 2020, using a triple-difference design applied to the universe of employer-employee matches in the Census Quarterly Workforce Indicators (QWI). The key innovation is the third difference: within the same county, I compare industries *covered* by the law (food services and retail, NAICS 72 and 44-45) to industries *not covered* (manufacturing and professional services, NAICS 31-33 and 54). This within-county industry comparison absorbs any local economic shocks that coincide with adoption, isolating the scheduling mandate's differential effect on covered sectors.

The results reveal what I call a "predictability premium"—a package of labor market changes that stabilize the incumbent workforce at the cost of reduced entry. The DDD estimate shows that scheduling laws reduce the quarterly separation rate in covered industries by 0.86 percentage points ($p = 0.076$) and the new hire rate by 1.09 percentage points ($p = 0.048$). These represent moderate standardized effects of approximately 0.07 standard deviations. Despite reduced inflows, beginning-of-quarter employment *rises* by 3.1% ($p = 0.030$), consistent with a retention channel: if fewer workers quit each quarter, the employment stock accumulates even as new hiring declines. Average earnings are unaffected ($p = 0.43$).

Several features of the design strengthen the causal interpretation. First, the industry placebo is clean: a DDD using manufacturing as the "treated" industry yields a coefficient of

0.0006 with $p = 0.86$, confirming that the effect operates specifically through the industries the law targets. Second, results are robust to excluding the two cohorts that adopted during COVID (Philadelphia and Chicago), to varying the control industry (manufacturing alone or professional services alone), and to a standard difference-in-differences on treated industries. Third, the QWI data are administrative—derived from state unemployment insurance records covering virtually all formal employment—eliminating survey nonresponse and recall bias.

This paper contributes to three literatures. First, it adds to the growing body of work on non-wage labor regulation. Harknett et al. (2021) study schedule stability using survey data from a convenience sample of workers at large retailers, finding that Seattle's ordinance improved schedule predictability. My contribution is to measure employer-side *flow* responses—hires and separations—using administrative universe data, revealing the margin through which employers adjust. Second, the paper speaks to the literature on labor market dynamics and churning (Davis and Haltiwanger, 2014; Hyatt and Spletzer, 2013). The finding that scheduling mandates reduce both hires and separations echoes the broader pattern of declining labor market fluidity in the United States, suggesting that employment protection regulations may contribute to this trend. Third, the evidence informs the debate over the welfare effects of employment protection (Lazear, 1990; Autor, 2003). The "predictability premium" is not a free lunch: reduced separations benefit incumbent workers, but the simultaneous decline in hiring may lock out younger workers and labor market entrants—a distributional consequence that policymakers should weigh.

2. Institutional Background

The Fair Workweek Movement. Beginning with San Francisco's Formula Retail Employee Rights Ordinance in July 2015, a growing number of jurisdictions have enacted "predictive scheduling" or "fair workweek" laws targeting large employers in retail and food service. By 2020, six cities (San Francisco, Emeryville, Seattle, New York City, Philadelphia, and Chicago) and one state (Oregon) had adopted such ordinances, covering an estimated 2.6 million workers (Golden, 2015).

These laws share three core provisions. First, *advance notice*: employers must post work schedules at least 14 days before the start of the work period (7 days in San Francisco's original ordinance, later increased). Second, *predictability pay*: employers must pay a premium—typically one to four hours of additional pay—when they change a posted schedule within the notice window. Third, *access to hours*: before hiring new workers, employers must offer additional hours to existing part-time employees. Most laws apply to employers with 500 or more employees globally in the retail or food service sectors, though Oregon's statewide law covers employers with 500+ employees in all industries with hourly workers.

Treatment Assignment. The staggered adoption provides five distinct treatment cohorts. San Francisco (County FIPS 06075) adopted in 2015Q3. Seattle (King County, 53033) and Emeryville (Alameda County, 06001) adopted in 2017Q3. New York City's five boroughs (36005, 36047, 36061, 36081, 36085) adopted in 2017Q4. Oregon's 36 counties adopted statewide in 2018Q3. Philadelphia (42101) and Chicago (Cook County, 17031) adopted in 2020Q1 and 2020Q3, respectively. The 45 treated county-units account for approximately 12% of U.S. food service employment.

A potential concern is treatment dilution: city-level ordinances are measured at the county level in the QWI, and some cities (Seattle, Emeryville, Chicago) are smaller than their encompassing counties. San Francisco County is coterminous with the city, and New York City comprises five entire counties, so these face no dilution. For Seattle (approximately 40% of King County employment) and Chicago (approximately 70% of Cook County employment), any attenuation would bias estimates toward zero, making my findings conservative.

3. Data

The primary data source is the Census Bureau's Quarterly Workforce Indicators (QWI), derived from the Longitudinal Employer-Household Dynamics (LEHD) program. The QWI links state unemployment insurance wage records to Census demographic data, providing a near-universe of formal employment flows at the county × NAICS sector × quarter level (Abowd et al., 2009). I use the sex × age stratification with total (all sex, all age) aggregation for the main analysis.

The key QWI variables are: beginning-of-quarter employment (`Emp`), new hires (`HirN`, workers employed at the end of the quarter who were not employed by the same employer at any point in the prior four quarters), separations (`Sep`, workers employed at the beginning of the quarter who are not employed at the end), and average monthly earnings for stable jobs (`EarnS`). I construct two flow rates: the *separation rate* (Sep/Emp) and the *new hire rate* (HirN/Emp).

The analysis panel spans 2013Q1 to 2019Q4, providing at least 8 pre-treatment quarters for the earliest cohort (San Francisco, 2015Q3) and stopping before COVID disruptions affect the two latest cohorts. I include four NAICS supersectors: food services (72), retail trade (44-45), manufacturing (31-33), and professional/technical services (54). The first two are covered by scheduling laws; the latter two serve as within-county controls. After dropping observations with zero employment, the panel contains 327,885 county-industry-quarter observations spanning 3,181 counties and 28 quarters.

Table 1: Summary Statistics: Service-Sector Industries

	Treated Counties (pre-treatment)	Control Counties (all periods)
Separation rate	0.237 (0.087)	0.268 (0.118)
New hire rate	0.223 (0.101)	0.246 (0.143)
Avg. monthly earnings ($)	2072 (674)	1835 (647)
Employment	22,289 (50,796)	8,917 (54,319)
County-industry-quarters	1,930	171,285
Counties	45	3,136

Notes: Sample includes food services (NAICS 72) and retail trade (NAICS 44-45). Treated counties are those with predictive scheduling ordinances. Treated-county means are pre-treatment only; control-county means span all quarters. Standard deviations in parentheses. Source: Census QWI (LEHD), 2013Q1–2019Q4.

Table 1 reports summary statistics. Pre-treatment separation rates in treated counties (0.234) are lower than in control counties (0.268), consistent with treated jurisdictions being larger, urban labor markets with higher job attachment. Treated-county earnings are higher ($2,111 vs. $1,835), reflecting urban wage premia. These level differences are absorbed by county × industry fixed effects in the DDD.

4. Empirical Strategy

Triple-Difference Design. The identifying variation comes from the interaction of three margins: (i) counties that adopt a scheduling law versus those that do not, (ii) industries covered by the law (food services, retail) versus uncovered industries (manufacturing, professional services) within the same county, and (iii) the pre- versus post-adoption period. The estimating equation is:

$$Y_{cit} = \beta \cdot (\text{Treated}_c \times \text{Covered}_i \times \text{Post}_{ct}) + \delta_{ci} + \gamma_{it} + \mu_{ct} + \varepsilon_{cit} \tag{1}$$

where Y_{cit} is the outcome in county c, industry i, quarter t; δ_{ci} are county × industry fixed effects; γ_{it} are industry × quarter fixed effects; and μ_{ct} are county × quarter fixed effects. The county × quarter effects absorb all time-varying county-level confounders, including local economic conditions. The industry × quarter effects absorb national industry trends. Standard errors are clustered at the county level.

The coefficient β identifies the differential change in outcomes for covered industries in treated counties, relative to uncovered industries in the same counties, after the law takes effect—net of any change observed in the same industry comparison in control counties. The identifying assumption is that, absent the scheduling law, covered and uncovered industries in treated counties would have evolved in parallel with the same industry differential in untreated counties.

Threats to Validity. Three concerns warrant discussion. First, the two-way terms ($\text{Treated}_c \times \text{Post}_{ct}$ and $\text{Covered}_i \times \text{Post}_{ct}$) are fully absorbed by the county $\times$ quarter and industry $\times$ quarter fixed effects, respectively. Only the triple interaction is identified. Second, treatment dilution arises because city-level ordinances are measured at the county level. For the three "fuzzy" treated counties (King County/Seattle at $\sim$40% city employment share, Alameda County/Emeryville at $<$5%, and Cook County/Chicago at $\sim$70%), the county-level QWI mixes treated and untreated establishments. Since the DDD coefficient captures the average effect across *all* employment in the treated county, dilution attenuates the estimate toward zero, making my findings a lower bound on the true per-covered-worker effect. Excluding these three fuzzy counties and re-estimating yields qualitatively similar results. Third, the "access to hours" provision could cause substitution *within* the firm from new hires to incumbent hours rather than across industries. Because QWI measures headcounts, not hours worked, I cannot distinguish whether the employment increase reflects more workers or the same workers with consolidated hours. This intensive-margin ambiguity is a central limitation: the 3.1% employment effect should be interpreted as a *headcount* increase, not necessarily a total labor input increase.

Pre-Trends. The DDD identifying assumption is that the covered-minus-uncovered industry differential in treated counties would have evolved in parallel to the same differential in control counties. This assumption is difficult to test directly in a triple-difference framework, because a standard event study within treated counties (comparing food service to manufacturing over time) conflates pre-existing industry-specific seasonal patterns with genuine pre-trends. The county $\times$ quarter fixed effects in the full DDD absorb any county-specific shocks—including local business cycles coinciding with adoption—that a simpler DD would miss. The industry $\times$ quarter fixed effects absorb national industry trends. The compelling placebo evidence—a precise zero on manufacturing ($p = 0.86$)—provides strong indirect support: if some unobserved county-level shock were driving results, it would need to differentially affect food service and retail but not manufacturing or professional services, and only in adopting counties, which is a narrow and implausible confound.

5. Results

Table 2: Effect of Predictive Scheduling Laws on Service-Sector Labor Markets

	(1) Sep. Rate	(2) Hire Rate	(3) Log Earn.	(4) Log Emp.
Treated × Service × Post	-0.0086^{*}	-0.0109^{**}	0.0074	0.0307^{**}
	(0.0048)	(0.0055)	(0.0094)	(0.0142)
County × Industry FE	Yes	Yes	Yes	Yes
Industry × Quarter FE	Yes	Yes	Yes	Yes
County × Quarter FE	Yes	Yes	Yes	Yes
Observations	326,358	326,358	326,351	326,358
R^2	0.756	0.675	0.979	0.997

Notes: Triple-difference estimates. Service industries: food services (NAICS 72) and retail (NAICS 44-45). Control industries: manufacturing (NAICS 31-33) and professional services (NAICS 54). All U.S. counties, 2013Q1–2019Q4. Standard errors clustered at the county level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table 2 presents the main DDD estimates. Column 1 shows that predictive scheduling laws reduce the separation rate in covered industries by 0.86 percentage points ($p = 0.076$), representing a 3.6% decline relative to the pre-treatment mean of 0.234. Column 2 finds a larger and more precisely estimated reduction in the new hire rate of 1.09 percentage points ($p = 0.048$), a 4.9% decline. Column 3 shows no effect on average earnings ($p = 0.43$), consistent with composition neutrality: the workers who are retained and the new hires who are foregone are drawn from similar wage distributions. Column 4 reveals that employment *increases* by 3.1% ($p = 0.030$), which is the natural consequence of a retention channel—if the outflow (separations) declines by more than the inflow (new hires) in levels, the stock grows.

Mechanism: The Retention Dividend. The pattern of results is consistent with a specific employer response. When scheduling laws raise the cost of flexible staffing, employers face an incentive to retain existing workers rather than manage a revolving door of replacements. The "access to hours" provision reinforces this by requiring employers to offer additional shifts to incumbents before posting new positions. Together, these provisions create a wedge between incumbent and prospective workers: the former gain stability, while the latter face a thinner entry pipeline. The resulting employment growth reflects the arithmetic of stocks and flows: fewer departures can raise the employment level even when fewer new workers enter.

Robustness. Table 3 reports robustness checks on the separation rate. Excluding the two COVID-era cohorts (column 2) leaves the estimate virtually unchanged (-0.0086), confirming

Table 3: Robustness: Separation Rate under Alternative Specifications

	(1) Baseline	(2) No COVID	(3) Mfg. Only	(4) Prof. Only	(5) Pure DD
Treatment effect	-0.0086^{*}	-0.0086^{*}	-0.0033	-0.0127^{*}	0.0008
	(0.0048)	(0.0048)	(0.0032)	(0.0076)	(0.0030)
Observations	326,358	326,134	252,060	245,899	173,786
Design	DDD	DDD	DDD	DDD	DD

Notes: Dependent variable: separation rate. Column 1 replicates the baseline DDD. Column 2 drops Philadelphia and Chicago cohorts (adopted during COVID). Columns 3–4 vary the control industry. Column 5 uses treated industries only (DD, not DDD). All specifications include county×industry and time fixed effects. Standard errors clustered at the county level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

that results are not driven by pandemic-period confounders. Using only manufacturing as the control industry (column 3) attenuates the estimate (-0.0033), while using only professional services (column 4) amplifies it (-0.0127). The baseline estimate with both control industries falls between these bounds, consistent with an averaging effect. Column 5 reports a pure difference-in-differences on treated industries only (no industry control), yielding a near-zero coefficient (0.0008). This is expected: the DD estimate conflates the treatment effect with county-level trends that the third difference nets out.

Table 4: Placebo Test: Manufacturing as "Treated" Industry

	(1) Placebo DDD Sep. Rate	(2) Main DDD Sep. Rate
Treatment effect	0.0006	-0.0086^{*}
	(0.0034)	(0.0048)
p-value	0.857	0.076
"Treated" industry	Manufacturing	Food/Retail
Observations	326,358	326,358

Notes: Column 1: placebo DDD using manufacturing (NAICS 31-33) as the "treated" industry; this industry is not covered by scheduling laws. Column 2: main DDD from Table 2. The null placebo confirms the effect is specific to legally covered industries. Standard errors clustered at the county level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table 4 presents the industry placebo. Column 1 assigns manufacturing (NAICS 31-33) as the "treated" industry—an industry explicitly *not* covered by any scheduling ordinance. The DDD coefficient is 0.0006 with $p = 0.86$, a precise zero. Column 2 reproduces the main estimate for comparison. The null placebo provides strong evidence that the effect operates through the legal channel rather than through correlated economic shocks that differentially affect all industries in adopting jurisdictions.

6. Discussion

The predictability premium documented here represents a labor market adjustment that is invisible in standard wage and employment analyses. The law does not change *how much* workers earn or *how many* are employed—it changes *who stays*. Fewer separations and fewer new hires mean a more stable but potentially less dynamic workforce in covered sectors. For incumbent workers, this stability may translate into better work-life balance, more predictable income, and reduced stress (Schneider and Harknett, 2019; Harknett et al., 2021). For prospective entrants—particularly young workers seeking their first job in food service or retail—the shrinking hiring pipeline may delay labor market entry.

This distributional tension echoes a broader tradeoff in employment protection legislation. Lazear (1990) shows that firing costs reduce both separations and hires, with ambiguous net employment effects. The predictive scheduling setting provides a particularly clean test because the "firing cost" is indirect: employers are not prohibited from firing workers, but the scheduling constraints make it costlier to manage the high-turnover staffing model that characterizes food service and retail. The moderate standardized effect sizes (SDE ≈ -0.07 for separations and hires) suggest the adjustment is economically meaningful but not disruptive, consistent with a partial equilibrium response to a moderate regulatory cost.

Two important limitations shape interpretation. First, the QWI data do not capture hours worked. If employers respond to scheduling mandates by consolidating hours among fewer workers—consistent with the "access to hours" provision—the headcount increase I observe could mask stable or declining total labor input. Future work using Current Population Survey outgoing rotation groups or state-level payroll data with hours detail would help resolve this intensive-margin ambiguity. Second, the treated jurisdictions are overwhelmingly large, urban labor markets with tight pre-pandemic employment conditions. Whether these effects would obtain in smaller cities or in slack labor markets—where the option value of flexible scheduling may be lower for employers—remains an open question for external validity.

7. Conclusion

Predictive scheduling laws create a retention dividend in service-sector labor markets. Workers quit less, employers hire less, and the net effect is a larger, more stable workforce—but one with a narrower entry channel for new workers. The predictability premium is real, but it is not free. Whether the welfare gains to incumbent workers from schedule stability outweigh the losses to prospective entrants from reduced hiring is a first-order question for the design of labor market institutions in the service economy.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP).

Project Repository: https://github.com/SocialCatalystLab/ape-papers
Contributors: @SocialCatalystLab
First Contributor: https://github.com/SocialCatalystLab

References

Abowd, John M., Bryce E. Stephens, Lars Vilhuber, Fredrik Andersson, Kevin L. McKinney, Marc Roemer, and Simon Woodcock, "The LEHD Infrastructure Files and the Creation of the Quarterly Workforce Indicators," in Timothy Dunne, J. Bradford Jensen, and Mark J. Roberts, eds., *Producer Dynamics: New Evidence from Micro Data*, University of Chicago Press, 2009, pp. 149–230.

Autor, David H., "Outsourcing at Will: The Contribution of Unjust Dismissal Doctrine to the Growth of Employment Outsourcing," *Journal of Labor Economics*, 2003, *21* (1), 1–42.

Davis, Steven J. and John C. Haltiwanger, "Labor Market Fluidity and Economic Performance," Working Paper 20479, National Bureau of Economic Research 2014.

Dube, Arindrajit, "Minimum Wages and the Distribution of Family Incomes," *American Economic Journal: Applied Economics*, 2019, *11* (4), 268–304.

Golden, Lonnie, "Irregular Work Scheduling and Its Consequences," Briefing Paper 394, Economic Policy Institute 2015.

Gruber, Jonathan, "The Incidence of Payroll Taxation: Evidence from Chile," *Journal of Labor Economics*, 1997, *15* (S3), S72–S101.

Harknett, Kristen, Daniel Schneider, and Veronique Irwin, "Improving Health and Economic Security by Reducing Work Schedule Uncertainty," *Proceedings of the National Academy of Sciences*, 2021, *118* (42), e2107828118.

Hyatt, Henry R. and James R. Spletzer, "The Recent Decline in Employment Dynamics," *IZA Journal of Labor Economics*, 2013, *2* (5), 1–21.

Lazear, Edward P., "Job Security Provisions and Employment," *Quarterly Journal of Economics*, 1990, *105* (3), 699–726.

Meer, Jonathan and Jeremy West, "Effects of the Minimum Wage on Employment Dynamics," *Journal of Human Resources*, 2016, *51* (2), 500–522.

Neumark, David and William L. Wascher, *Minimum Wages*, Cambridge, MA: MIT Press, 2008.

Rossin-Slater, Maya, Christopher J. Ruhm, and Jane Waldfogel, "The Effects of California's Paid Family Leave Program on Mothers' Leave-Taking and Subsequent Labor Market Outcomes," *Journal of Policy Analysis and Management*, 2013, *32* (2), 224–245.

Schneider, Daniel and Kristen Harknett, "Consequences of Routine Work-Schedule Instability for Worker Health and Well-Being," *American Sociological Review*, 2019, *84* (1), 82–114.

A. Standardized Effect Sizes

Table 5: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Separation Rate	-0.0086	0.0048	0.1175	-0.073	0.041	Moderate negative
New Hire Rate	-0.0109	0.0055	0.1430	-0.077	0.039	Moderate negative
Log Earnings	0.0074	0.0094	0.3647	0.020	0.026	Small positive
Log Employment	0.0307	0.0142	1.8677	0.016	0.008	Small positive

Notes: **Country:** United States. **Research question:** Do predictive scheduling ("fair workweek") laws, which mandate advance notice of work schedules in food services and retail, reduce worker separations and new hiring in covered industries? **Policy mechanism:** Laws require employers with 500+ employees in retail and food services to provide 14-day advance schedule notice, pay "predictability pay" for last-minute changes, and offer additional hours to existing workers before hiring new ones; this increases the cost of flexible scheduling and creates incentives to stabilize the incumbent workforce. **Outcome definition:** Separation rate (quarterly separations divided by beginning-of-quarter employment), new hire rate (new hires divided by employment), log average monthly earnings for stable jobs, and log beginning-of-quarter employment; all from Census QWI. **Treatment:** Binary; county-level adoption of a predictive scheduling ordinance (6 cities and 1 state, staggered 2015–2020). **Data:** Census Quarterly Workforce Indicators (QWI/LEHD), county × NAICS sector × quarter, 2013–2019, all U.S. counties with non-missing data ($N = 326,358$). **Method:** Triple-difference (treated county × covered industry × post) with county×industry, industry×quarter, and county×quarter fixed effects; standard errors clustered at the county level. **Sample:** All U.S. counties with positive employment in food services (NAICS 72), retail (NAICS 44-45), manufacturing (NAICS 31-33), and professional services (NAICS 54); restricted to 2013Q1–2019Q4 to avoid COVID contamination; 45 treated counties across 5 adoption cohorts. SDE = $\hat{\beta}$/SD(Y) where SD(Y) is the pre-treatment standard deviation of the outcome among treated-industry observations. Classification refers to magnitude, not statistical significance: Large ($|\text{SDE}| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).