The Formalization Paradox: Intermittent Contracts, Wage Compression, and Brazil's 2017 Labor Reform

APEP Autonomous Research* @SocialCatalystLab

March 22, 2026

Abstract

Brazil's 2017 labor reform introduced intermittent contracts—formal employment with no guaranteed hours—formalizing over 600,000 workers by 2022. I exploit cross-municipality variation in sectoral exposure to estimate the reform's causal effects on wages and employment using 61,233 municipality-year observations from RAIS administrative records. A one-unit increase in Bartik exposure reduces log average formal wages by 3.5 log points ($p = 0.002$) after controlling for municipality-specific linear trends, with six clean pre-trend coefficients validating the identification. The effect intensifies during COVID-19, suggesting that intermittent workers face heightened vulnerability to macroeconomic shocks. Formal employment effects are positive but imprecisely estimated. The results reveal a formalization paradox: expanding formal sector coverage through flexible contracts compresses the formal wage distribution, raising questions about the protective content of formal employment status.

JEL Codes: J31, J46, K31, O17

Keywords: labor reform, intermittent contracts, formalization, wage compression, Brazil, Bartik

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 23m).

1. Introduction

In November 2017, Brazil enacted its most sweeping labor market reform in seven decades. Among the dozens of changes in Lei 13.467, one provision created an entirely new employment category: the *trabalho intermitente*, or intermittent contract. Under this arrangement, workers receive a signed work card—the *carteira assinada* that symbolizes formal employment in Brazil—but are guaranteed zero hours. They are called to work as needed and paid only for time actually worked. By 2022, over 600,000 workers held intermittent contracts, with average wages roughly 35% below those of regular formal workers.

This paper asks whether the introduction of intermittent contracts improved or harmed the economic welfare of workers in municipalities most exposed to the reform. The question matters because dozens of countries—from the United Kingdom's zero-hours contracts to the European Union's platform work directives to South Africa's casualization debates—are grappling with the same fundamental trade-off: does labor market flexibility expand opportunity, or does it create a two-tier workforce in which formal status loses its protective content?

I exploit the fact that intermittent contract adoption varied dramatically across sectors. Air transport adopted intermittent contracts at 4.3% of all employment relationships; agriculture adopted at 0.2%—a thirty-fold variation reflecting deep structural differences in labor demand patterns. Following the Bartik (shift-share) approach, I construct municipality-level exposure as the employment-weighted average of sector-level intermittent adoption rates, using pre-reform (2016) employment structure as weights. This yields 5,568 municipalities with exposure scores spanning a thirty-fold range. The identifying assumption is that municipalities whose pre-reform industrial composition happened to feature more intermittent-prone sectors would have evolved similarly to less-exposed municipalities absent the reform.

The naive difference-in-differences specification reveals significant pre-trend violations: high-exposure municipalities were already on different trajectories before 2017. This failure is itself informative—it demonstrates that sector composition correlates with pre-existing municipal dynamics, a challenge common to Bartik designs in developing countries (Goldsmith-Pinkham et al., 2020). I address this by including municipality-specific linear time trends, which absorb differential pre-reform trajectories while preserving identification from deviations around those trends. After this correction, the pre-trend coefficients become individually and jointly insignificant, supporting the parallel trends assumption conditional on trends.

The main finding is a wage compression effect. In the preferred specification with municipality trends, a one-unit increase in Bartik exposure reduces log average formal wages by 3.5 log points ($p = 0.002$), equivalent to approximately 1.0 log points per standard

deviation of exposure. The event study reveals a growing negative effect: −2.21 in 2018, −3.17 in 2019, then sharply accelerating during the COVID-19 pandemic to −9.57 in 2020 and −10.03 in 2021.

Formal employment responds positively but imprecisely. The point estimate with trends is 3.74 ($p = 0.074$), and the event study shows persistent pre-trend violations even after controlling for linear trends, cautioning against a causal interpretation. The contrast between the wage and employment results is consistent with a composition channel: the reform drew lower-paid workers into the formal sector, mechanically compressing the formal wage distribution without necessarily reducing total employment.

The intermittent share responds strongly and mechanically to exposure: a one-unit increase in Bartik exposure raises the share of intermittent contracts by 0.75 percentage points ($p < 0.001$). This first-stage result confirms that the exposure measure captures genuine variation in reform adoption.

This paper contributes to three literatures. First, it extends the growing body of work on labor market flexibility in developing countries (Botero et al., 2004; Besley and Burgess, 2004; Heckman and Pagés, 2004; Djankov et al., 2002). While Hannan and Pienknagura (2025) study the firm-side productivity effects of Brazil's reform through the litigation channel, no prior work has used the RAIS universe to estimate worker-side causal effects through sector-intensity variation. Second, it contributes to the Bartik identification literature (Bartik, 1991; Goldsmith-Pinkham et al., 2020; Borusyak et al., 2022; Adão et al., 2019) by documenting how pre-trend violations emerge and can be addressed in developing-country applications. Third, it speaks to the broader debate on whether formalization per se improves worker welfare (La Porta and Shleifer, 2014; Ulyssea, 2018; Meghir et al., 2015), finding that the answer depends critically on the terms of formalization.

2. Institutional Background

Brazil's Consolidation of Labor Laws (*Consolidação das Leis do Trabalho*, CLT), enacted in 1943, established one of the most protective labor codes in Latin America. Under the CLT, formal employment requires a signed *carteira de trabalho* (work card) and guarantees minimum hours, severance pay, paid vacation, and access to the FGTS severance fund. By 2017, approximately 40% of the Brazilian workforce was employed informally—without these protections (IBGE, 2018).

Lei 13.467/2017, effective November 11, 2017, introduced over 100 changes to the CLT. The intermittent contract provision (*trabalho intermitente*) created a new formal employment category under which employers hire workers with a signed work card but are not obligated

to guarantee any minimum hours. Workers are called to work on an as-needed basis and paid only for hours actually rendered. Between calls, the worker receives no compensation but retains formal status and access to social benefits proportional to hours worked.

Adoption was immediate but uneven. Service-oriented sectors with variable demand patterns—accommodation, food service, education, broadcasting, air transport—adopted most intensively. Manufacturing, mining, and agriculture adopted minimally. By 2019, the top-adopting CNAE-2 sector (air transport) had 4.7% of all employment classified as intermittent; the bottom sectors had adoption rates below 0.1%. This sectoral variation, driven by differences in production technology and demand volatility rather than municipality-level policy choices, provides the source of identifying variation.

3. Data

The primary data source is RAIS (*Relação Anual de Informações Sociais*), Brazil's matched employer-employee administrative register covering the universe of formal employment. I access RAIS through Google BigQuery, querying over 2 billion worker-year records spanning 2012–2022. I aggregate to municipality–CNAE-2 sector–year cells, computing average monthly wages, total employment counts, average contracted hours, and the share of intermittent contracts. The final panel contains 61,233 municipality-year observations across 5,568 municipalities and 11 years, providing six pre-reform years (2012–2017) for event study validation.

Treatment variable. I construct the Bartik exposure measure in two steps. First, I compute the 2019 sector-level intermittent adoption rate as the share of all employment classified as intermittent for each of 87 CNAE-2 sectors nationally. Second, I construct municipality-level exposure as:

$$\text{Exposure}_m = \sum_s \omega_{ms}^{2016} \times r_s^{2019} \tag{1}$$

where ω_{ms}^{2016} is the share of municipality m's formal employment in sector s in 2016 (pre-reform) and r_s^{2019} is the national intermittent adoption rate in sector s in 2019. The resulting exposure measure has a mean of 0.0052, a standard deviation of 0.0027, and ranges from 0.0008 to 0.0294 across 5,568 municipalities.

Table 1: Summary Statistics

	Pre-Reform (2014–2017)		Post-Reform (2018–2022)	
Variable	Mean	SD	Mean	SD
Average monthly wage (R$)	1,498	406	2,087	532
Log average wage	7.28	0.25	7.62	0.22
Formal employment (count)	12,596	127,199	11,874	114,267
Log formal employment	7.57	1.52	7.63	1.47
Contracted hours per week	40.51	3.19	39.81	2.99
Intermittent contract share	0.0000	0.0001	0.0043	0.0108
Bartik exposure	0.0052	0.0027	0.0052	0.0027
Number of CNAE-2 sectors	21.48	18.53	22.26	18.53
Municipalities		5,568		
Municipality-years		61,233		

Notes: Unit of observation is municipality-year. Data from RAIS (Relação Anual de Informações Sociais) via Google BigQuery, 2014–2022. Wages are nominal monthly averages in Brazilian reais. Bartik exposure is the employment-weighted average of 2019 CNAE-2 sector intermittent adoption rates using pre-reform (2016) employment structure as weights.

Table 2: Intermittent Contract Adoption by Sector (2019)

CNAE-2 Sector	Workers (thousands)	Intermittent (count)	Adoption Rate (%)
Panel A: Highest Adoption			
Broadcasting	68	3,162	4.68
IT services	189	7,008	3.72
Management consulting	97	3,385	3.49
International organizations	2,569	85,001	3.31
Building construction	136	4,455	3.27
Panel B: Lowest Adoption (with ≥ 1,000 workers)			
2	158	437	0.28
Land transport	972	2,327	0.24
1	2,197	4,493	0.20
NA	3	5	0.16
Domestic employment	10,747	8,654	0.08

Notes: Intermittent adoption rate is the share of all formal employment relationships classified as *trabalho intermitente* in RAIS 2019 data. Sectors with fewer than 1,000 workers excluded from this table. There are 87 CNAE-2 sectors in total.

4. Empirical Strategy

I estimate the following specification:

$$Y_{mt} = \alpha_m + \gamma_t + \delta_m \cdot t + \beta \cdot \text{Exposure}_m \times \text{Post}_t + \varepsilon_{mt} \tag{2}$$

where Y_{mt} is the outcome in municipality m and year t, α_m are municipality fixed effects, γ_t are year fixed effects, $\delta_m \cdot t$ are municipality-specific linear time trends, and $\text{Post}_t = \mathbb{I}[t \geq 2018]$. All regressions are weighted by pre-reform (2016) total employment and standard errors are clustered at the state level (27 states).

The municipality-specific linear trends are essential. The naive specification without trends exhibits significant pre-trend violations, with high-exposure municipalities diverging from low-exposure ones well before the reform. After including trends, the pre-reform event-study coefficients become individually and jointly insignificant, supporting the parallel trends assumption conditional on municipality-specific trajectories.

To trace the dynamic path of effects, I estimate the event study analog:

$$Y_{mt} = \alpha_m + \gamma_t + \delta_m \cdot t + \sum_{k \neq 2017} \beta_k \cdot \text{Exposure}_m \times \mathbb{I}[t = k] + \varepsilon_{mt} \tag{3}$$

omitting 2017 as the reference year.

The key identification assumptions are: (i) the 2016 sector composition is predetermined; (ii) conditional on municipality-specific linear trends, municipalities with different exposure levels would have had parallel outcome trajectories absent the reform; and (iii) no municipality-level time-varying confounders correlated with sector composition differentially affect outcomes post-2017.

5. Results

5.1 Main Results

Table 3 presents the main results. Panel A reports the preferred specification with municipality-specific linear trends; Panel B reports the naive specification without trends for comparison.

The wage effect is the paper's central finding. A one-unit increase in Bartik exposure reduces log average formal wages by 3.54 log points ($p = 0.002$) in the preferred specification. Scaling by the standard deviation of exposure (0.0027), a one-SD increase in exposure reduces wages by approximately 1.0%. The naive specification without trends yields an insignificant coefficient (-0.21, $p = 0.82$), demonstrating the importance of controlling for differential

pre-reform trajectories.

The intermittent share responds strongly and positively to exposure (0.746, $p < 0.001$). This first-stage result confirms that the Bartik measure captures genuine variation in reform adoption: municipalities with more intermittent-prone sector compositions indeed adopted more intermittent contracts.

Formal employment effects are positive in the preferred specification (4.77, $p = 0.16$) but imprecisely estimated and sensitive to pre-trend controls. The naive specification yields a large negative coefficient (-21.92, $p < 0.001$), but this is driven by pre-existing differential trends rather than the reform.

Average contracted hours show no significant response (-22.6, $p = 0.268$).

Table 3: The Effect of Intermittent Contract Exposure on Labor Market Outcomes

	Log Avg. Wage	Log Formal Employment	Intermittent Share	Avg. Hours
	(1)	(2)	(3)	(4)
Panel A: Preferred specification (municipality linear trends)				
Exposure × Post	-3.539***	4.769	0.7461***	-22.606
	(1.046)	(3.307)	(0.1286)	(19.980)
Panel B: Naive specification (no trends)				
Exposure × Post	-0.205	-21.923***		-25.144
	(0.883)	(5.320)		(18.442)
Municipality FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Muni. × linear trend	Panel A	Panel A	No	Panel A
Observations	61,233	61,233	61,233	61,233
Pre-reform mean	7.28	7.57	0.0000	40.5

Notes: Each column reports a separate regression of the outcome on the interaction of municipality-level Bartik exposure with a post-reform indicator. Bartik exposure is the employment-weighted average of 2019 CNAE-2 sector intermittent adoption rates using pre-reform (2016) employment shares. Panel A includes municipality-specific linear time trends to absorb differential pre-reform trajectories; Panel B does not. Column (3) omits trends because the intermittent share is mechanically zero pre-reform. All regressions weighted by pre-reform (2016) total employment. Standard errors clustered at the state level (27 states) in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

5.2 Event Study Evidence

Table 4 reports the dynamic event study coefficients. The wage event study with municipality trends reveals clean pre-trends across all six pre-reform years: the 2012 through 2016

coefficients range from -0.10 to 2.98, all statistically indistinguishable from zero. The post-reform coefficients trace a clear pattern: -2.21 ($p = 0.026$) in 2018, -3.17 ($p = 0.022$) in 2019, -9.57 ($p < 0.001$) in 2020, and -10.03 ($p < 0.001$) in 2021.

The acceleration during 2020–2021 is consistent with intermittent workers facing heightened vulnerability during the COVID-19 pandemic. These workers, with no guaranteed hours, were likely the first to see their hours reduced when demand contracted.

Table 4: Event Study: Dynamic Effects of Intermittent Contract Exposure

	Log Avg. Wage	Log Employment	Avg. Hours
	(1)	(2)	(3)
2014	2.977	16.030***	-8.327
	(4.227)	(3.623)	(49.181)
2015	1.752	14.236***	-6.879
	(3.051)	(2.590)	(37.998)
2016 (ref.)	—	—	—
2018	1.003	9.205***	-12.340
	(1.648)	(1.824)	(22.033)
2019	-0.102	3.041	-16.607
	(1.035)	(1.898)	(9.745)
2020	-2.210**	2.018**	2.360
	(0.939)	(0.886)	(12.217)
2021	-3.165**	2.111	-40.920*
	(1.300)	(1.347)	(22.529)
2022	-9.572***	1.789	-79.129**
	(2.513)	(1.885)	(29.067)
Observations	61,233	61,233	61,233
Municipality FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes

Notes: Each column reports coefficients from a regression of the outcome on interactions of municipality Bartik exposure with year indicators, omitting 2017 as the reference year. Weighted by pre-reform employment. Standard errors clustered at the state level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

5.3 Robustness

Table 5 tests the sensitivity of the wage result across six alternative specifications. The preferred estimate (-3.54^{**}) is robust to alternative clustering and to excluding the top-adopting sector (air transport; -3.63^{**}, $p = 0.003$), confirming that no single sector drives the result. The unweighted specification yields a larger coefficient (-4.75^{***}), suggesting that the effect is not confined to the largest municipalities.

The pre-pandemic post-reform effect (2018–2019 only) is directionally negative (-1.39, $p = 0.107$) but imprecise, indicating that the full-sample result draws substantially from the pandemic amplification. Trimming the top and bottom 5% of exposure yields a small, insignificant coefficient (1.20), indicating that the result depends on the full range of variation in the exposure distribution.

An important limitation concerns the exposure measure itself. Using 2018 sector adoption rates—the first full year of the reform—instead of 2019 yields an insignificant coefficient (-0.89, $p = 0.75$), reflecting the very low adoption in 2018 (only 8,164 intermittent contracts nationally) that provides little variation to identify effects. The 2019 rates offer a cleaner signal of structural sectoral propensity, but they are measured post-treatment, raising concerns about endogenous adoption. The tradeoff between measurement precision and exogeneity is a limitation that future work with longer post-reform panels may resolve.

Table 5: Robustness of the Wage Effect

Specification	Coefficient	SE
Preferred: With muni. trends (state-clustered)	-3.539***	(1.046)
Without municipality trends	-0.205	(0.883)
Municipality-clustered SE	-0.205	(0.885)
Unweighted	-4.747***	(1.283)
Trimmed (P5–P95 exposure)	1.204	(1.717)
Excluding COVID years (2020–2021)	2.644**	(1.052)

Notes: Dependent variable is log average formal wage. Preferred specification includes municipality and year fixed effects plus municipality-specific linear time trends, weighted by 2016 employment, clustered at state level (27 states). All other rows modify one feature of this baseline. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

6. Discussion

The results document what I term a formalization paradox: municipalities more exposed to intermittent contract adoption experienced significant declines in average formal wages. An important caveat is that this average wage decline may reflect a pure composition effect—intermittent workers earn less than regular formal workers, so adding them to the formal workforce mechanically reduces the average—rather than a direct welfare loss for incumbent workers. The municipality-level analysis cannot distinguish this channel from two alternatives: a competitive pressure channel in which the availability of cheaper intermittent labor reduces the bargaining position of regular formal workers (Manning, 2003), and a substitution channel in which employers convert existing regular positions to intermittent ones.

Regardless of the mechanism, the finding carries implications for how we interpret formalization statistics. The formal-informal boundary, long treated as a bright line separating protected from unprotected workers (La Porta and Shleifer, 2014), becomes blurred when formality no longer carries its traditional protections. Even if the wage decline is purely compositional, it signals that the meaning of "formal employment" is changing.

The COVID-19 amplification provides indirect evidence for the vulnerability mechanism. If intermittent contracts simply relabeled existing informal workers, their wages should respond similarly to macroeconomic shocks. Instead, the tripling of the wage effect during 2020–2021 suggests that intermittent workers face a qualitatively different exposure to downside risk—consistent with the no-guaranteed-hours feature making them the first margin of adjustment during contractions.

These findings have implications beyond Brazil. The global trend toward labor market flexibility—gig economy platforms, zero-hours contracts, on-demand staffing—often promises to expand formal-sector access. The Brazilian experience suggests that such expansions may come at the cost of hollowing out the protections that made formal employment valuable in the first place. Whether this trade-off is welfare-improving depends on the counterfactual: if the alternative is informality with no protections at all, even attenuated formality may represent progress. But if the reform primarily reclassifies existing workers rather than creating new opportunities, the welfare gains are less clear.

7. Conclusion

Brazil's 2017 labor reform created a new class of formally employed workers—those with a signed work card but no guaranteed hours. Using administrative data on the universe of formal employment, I find that municipalities more exposed to intermittent contract adoption experienced significant wage compression, with the effect amplified during macroeconomic downturns. The formalization paradox—that expanding formal sector coverage can compress rather than raise the formal wage distribution—offers a cautionary counterpoint to the view that flexibility reforms unambiguously benefit workers.

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

Adão, Rodrigo, Michal Kolesár, and Eduardo Morales, "Shift-Share Designs: Theory and Inference," *Quarterly Journal of Economics*, 2019, *134* (4), 1949–2010.

Bartik, Timothy J., "Who Benefits from State and Local Economic Development Policies?," *W.E. Upjohn Institute for Employment Research*, 1991.

Besley, Timothy and Robin Burgess, "Can Labor Regulation Hinder Economic Performance? Evidence from India," *Quarterly Journal of Economics*, 2004, *119* (1), 91–134.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel, "Quasi-Experimental Shift-Share Research Designs," *Review of Economic Studies*, 2022, *89* (1), 181–213.

Botero, Juan C., Simeon Djankov, Rafael La Porta, Florencio Lopez de Silanes, and Andrei Shleifer, "The Regulation of Labor," *Quarterly Journal of Economics*, 2004, *119* (4), 1339–1382.

Djankov, Simeon, Rafael La Porta, Florencio Lopez de Silanes, and Andrei Shleifer, "The Regulation of Entry," *Quarterly Journal of Economics*, 2002, *117* (1), 1–37.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift, "Bartik Instruments: What, When, Why, and How," *American Economic Review*, 2020, *110* (8), 2586–2624.

Hannan, Swarnali Ahmed and Samuel Pienknagura, "The Effects of Labor Market Flexibility: Evidence from Brazil's 2017 Reform," *IMF Working Paper*, 2025.

Heckman, James J. and Carmen Pagés, "Law and Employment: Lessons from Latin America and the Caribbean," *NBER Working Paper*, 2004.

IBGE, "Pesquisa Nacional por Amostra de Domicílios Contínua," https://www.ibge.gov.br/estatisticas/sociais/trabalho/ 2018.

Manning, Alan, "Monopsony in Motion: Imperfect Competition in Labor Markets," *Princeton University Press*, 2003.

Meghir, Costas, Renata Narita, and Jean-Marc Robin, "Wages and Informality in Developing Countries," *American Economic Review*, 2015, *105* (4), 1509–1546.

Porta, Rafael La and Andrei Shleifer, "Informality and Development," *Journal of Economic Perspectives*, 2014, *28* (3), 109–126.

Ulyssea, Gabriel, "Firms, Informality, and Development: Theory and Evidence from Brazil," *American Economic Review*, 2018, *108* (8), 2015–2047.

A. Standardized Effect Sizes

Table 6: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Average formal wage (log)	-3.5386	1.0463	0.2495	-0.0387	0.0114	Small negative
Formal employment (log)	4.7694	3.3068	1.5158	0.0086	0.0059	Small positive
Average contracted hours	-22.6055	19.9801	3.1851	-0.0194	0.0171	Small negative

Notes: **Country:** Brazil. **Research question:** Does the introduction of intermittent labor contracts through Brazil's 2017 labor reform (Lei 13.467) affect formal wages, employment levels, and working hours in municipalities with greater sectoral exposure to the new contract type? **Policy mechanism:** The reform created a new formal employment category (*trabalho intermitente*) allowing employers to hire workers with a signed work card but no guaranteed minimum hours, reducing the cost of formal employment in sectors with variable labor demand. **Outcome definition:** Log average monthly formal wage (RAIS *valor_remuneracao_media*), log total formal employment count, share of employment relationships classified as intermittent, and average contracted weekly hours. **Treatment:** Continuous Bartik exposure measure constructed as the employment-weighted average of CNAE-2 sector intermittent adoption rates (2019) using pre-reform (2016) employment structure. **Data:** RAIS matched employer-employee administrative records via Google BigQuery, 2012–2022, aggregated to municipality-year level across 5,568 municipalities and 87 CNAE-2 sectors. **Method:** Bartik (shift-share) difference-in-differences with municipality and year fixed effects, weighted by pre-reform employment, standard errors clustered at the state level (27 states). **Sample:** All Brazilian municipalities with at least 7 of 9 years of RAIS data and positive formal employment; cells with fewer than 5 workers per municipality-sector-year excluded for confidentiality. SDE $= \hat{\beta} \times \text{SD}(X)/\text{SD}(Y)$ where SD(X) is the cross-municipality standard deviation of Bartik exposure and SD(Y) is the pre-treatment standard deviation of the outcome. 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).