The Cost of Sponsorship: Kafala Reform, Monopsony, and Firm Value in the UAE

APEP Working Paper

March 2026

Abstract

The kafala system ties migrant workers to their employers, creating textbook monopsony. I study the UAE's 2022 abolition of the employer no-objection-certificate requirement using a cross-sectional event study on 45 Dubai Financial Market firms. Comparing high-exposure sectors (real estate, services, industrial) to low-exposure sectors (banking, insurance, telecom) across three legislative events, I find a precisely estimated null: high-exposure firms earned 3.59 percentage points higher cumulative abnormal returns than low-exposure firms ($p = 0.387$; randomization inference $p = 0.354$). Zero of five placebo dates are significant, and GCC benchmark firms show no reaction. Under standard event-study assumptions, the 95 percent confidence interval bounds the differential valuation effect at less than 4.5 percent of listed firm value. Concurrent Emiratisation quotas complicate the interpretation—the null may reflect offsetting cost shocks rather than the absence of monopsony rents. The result constrains the range of plausible rent magnitudes but does not isolate kafala rents from the broader reform package.

JEL Codes: J42, J61, G14, O15, F22

Keywords: Kafala, monopsony, migrant labor, event study, UAE, labor market reform, worker mobility

1 Introduction

In the United Arab Emirates, nine out of ten private-sector workers cannot change jobs without their employer's written permission. This institutional lock—known as the *kafala* system—ties a migrant worker's legal residency to a specific employer, the "sponsor," who must approve any job change and who can cancel the worker's visa at will. The system governs roughly 35 million workers across the Gulf Cooperation Council states. To human rights organizations, it is modern servitude. To economists, it is something more precise: textbook monopsony.

The monopsony framework (Robinson, 1933; Manning, 2003) makes a sharp prediction about kafala. By restricting worker mobility, the system depresses wages below marginal revenue product, transferring rents from workers to firms. The employer's no-objection certificate (NOC)—the requirement that workers obtain their current employer's permission before switching jobs—is the institutional mechanism that generates this markdown. Remove the NOC, and monopsony power falls. Workers can now credibly threaten to leave. Wages should rise toward competitive levels. Firm profits should decline.

On September 20, 2021, the UAE signed Federal Decree-Law No. 33, which among other provisions abolished the NOC requirement for private-sector workers. Implementing regulations followed on November 15, 2021, and the law took effect on February 2, 2022. This was the most sweeping kafala reform in any GCC state—broader in scope than Saudi Arabia's partial mobility reforms and more comprehensive than Qatar's 2020 changes (International Labour Organization, 2020). The reform provides a rare natural experiment: a sudden, large reduction in employer monopsony power over a predominantly migrant labor force.

This paper asks how much firms benefited from kafala's restriction on worker mobility. The estimand is the unanticipated change in expected discounted profits for high- versus low-exposure listed firms, attributable to information revealed at the three legislative milestones, relative to what was expected just before each milestone. I answer this question using stock

market data. If the kafala system generated economically significant monopsony rents, then the reform's announcement should have reduced the expected future profits of firms that relied most heavily on kafala-bound workers. Real estate developers, service companies, and industrial firms—where migrant workers constitute 85 to 95 percent of the workforce—should have experienced negative abnormal returns relative to banks, insurers, and telecom firms, where the workforce was already relatively mobile and the reform's bite was weaker.

I implement a cross-sectional event study on 45 firms listed on the Dubai Financial Market (DFM), spanning January 2019 through December 2024. I classify 18 firms in real estate, services, and industrial sectors as "high exposure" to kafala reform, and 27 firms in banking, insurance, telecom, financial services, and investment as "low exposure." I estimate cumulative abnormal returns (CARs) around each of three reform-related events and test whether high-exposure firms experienced differentially negative returns.

The main result is a well-identified null. In the preferred specification using a $[-1, +3]$ event window pooled across all three events, high-exposure firms earned 3.59 percentage points *higher* CARs than low-exposure firms, with a standard error of 4.11 and a p-value of 0.387. A continuous exposure measure yields a coefficient of 16.6 percentage points per unit of migrant share ($p = 0.320$). A market-model specification produces a coefficient of 5.19 percentage points ($p = 0.195$). The stacked difference-in-differences estimator yields a near-zero coefficient. Randomization inference, which permutes treatment assignment 1,000 times, gives a p-value of 0.354. None of these estimates is statistically distinguishable from zero, and—importantly—none is negative.

The null result is robust. I run the identical specification on five placebo dates distributed across the pre- and post-reform period; zero of five produce a significant coefficient. I show that labor-intensive firms on GCC benchmark exchanges (Qatar National Bank, Saudi Al Rajhi Bank, Saudi Aramco) did not react to UAE reform dates, confirming that the events are UAE-specific. Alternative event windows ranging from $[-1, +1]$ to $[0, +10]$ produce coefficients between 0.04 and 3.93 percentage points, all statistically insignificant. Leave-

one-out analysis confirms that no single firm or sector drives the result. The finding is a precisely estimated zero, not a failure of statistical power: the 95 percent confidence interval on the main coefficient is $[-4.5\%, +11.6\%]$, which rules out negative effects larger than 4.5 percentage points of firm value.

This paper makes three contributions. First, it provides the first market-based estimate of the differential valuation effect of kafala reform on labor-intensive firms. Under the stated assumptions—no anticipation, correct exposure classification, and no confounding events—the 95 percent confidence interval bounds the net valuation effect at less than 4.5 percent of listed firm value. This is an informative bound: if kafala transferred 10 or 20 percent of firm value from workers—as the system's extreme restrictions might suggest—we would have detected it. However, because the reform was bundled with Emiratisation quota intensification, the estimand is the net effect of the reform package, not kafala rents in isolation.

Second, the paper connects the Gulf labor literature to the broader monopsony research agenda. A growing body of work documents employer wage-setting power in developed-country labor markets (Card et al., 2018; Azar et al., 2022; Benmelech et al., 2022; Prager and Schmitt, 2021; Bassier et al., 2022); a recent meta-analysis finds a median markdown of approximately 15 percent (Sokolova and Sorensen, 2021). The estimated monopsony markdowns in these settings are typically 10 to 25 percent of wages. The kafala system—with its near-total restriction on worker mobility—should produce markdowns at the extreme upper end of this distribution. That the stock market placed a value of essentially zero on these rents is a puzzle that demands explanation. I discuss four interpretations: market anticipation, limited reform bite, bundling with offsetting Emiratisation costs, and the gap between de jure and de facto enforcement.

Third, the paper contributes to the event study methodology for policy evaluation in thin markets. With 45 firms and substantial cross-sectional variation in returns, I show that randomization inference and stacked multi-event designs can provide credible inference even in settings where asymptotic cluster-robust standard errors may be unreliable. The

comprehensive robustness battery—five placebo dates, four event windows, GCC placebo, market model, continuous treatment, leave-one-out—establishes a template for event studies on small exchanges.

2 Background

2.1 The Kafala System

The kafala (sponsorship) system is the dominant labor market institution governing migrant employment in the GCC states: Bahrain, Kuwait, Oman, Qatar, Saudi Arabia, and the UAE. Under kafala, every foreign worker's residency permit is tied to a specific employer, known as the *kafeel* (sponsor). The sponsor applies for the worker's entry visa, bears nominal financial responsibility, and—critically—retains the legal authority to approve or deny the worker's ability to change employers, leave the country, or in some cases move between emirates within a country.

The system's defining feature is the No Objection Certificate (NOC). Under pre-reform UAE law, a worker who wished to change employers had to obtain written consent from their current sponsor. If the sponsor refused, the worker faced two options: remain in the current job or leave the country, typically with a labor ban preventing re-entry for six to twelve months. This asymmetry in outside options is the source of the system's monopsony power. The worker's reservation wage is not determined by the next-best job offer in the UAE labor market—it is determined by the much lower wages available in their country of origin, minus the costs of return migration and future re-entry barriers.

The scale of kafala is difficult to overstate. The UAE's private-sector workforce is approximately 88 percent non-national (International Labour Organization, 2024). In construction, the share exceeds 95 percent; in services and hospitality, it exceeds 90 percent (Ministry of Human Resources and Emiratisation, 2023). Approximately 7 million migrant workers in the UAE, predominantly from South Asia and Southeast Asia, were governed by kafala

provisions as of 2021. Across all GCC states, the figure reaches roughly 35 million.

Human rights organizations have extensively documented the system's abuses: passport confiscation, wage theft, unsafe working conditions, and debt bondage through recruitment fees. From an economic perspective, these abuses are symptoms of a structural feature: kafala generates employer monopsony power by restricting workers' outside options. The question this paper asks is not whether kafala harms workers—the qualitative evidence is overwhelming—but whether the rents it generates for employers are large enough to be capitalized into firm values.

2.2 The 2022 Reform: Federal Decree-Law No. 33

The UAE's reform of its labor law was announced in stages over approximately five months:

1. **September 20, 2021:** The UAE government signed Federal Decree-Law No. 33 of 2021 Concerning the Regulation of Employment Relationships (UAE Government, 2021). The law's key provisions included: abolition of the NOC requirement, allowing workers to change employers without sponsor consent; introduction of fixed-term employment contracts (replacing the previous unlimited-contract system); provisions for part-time and flexible work arrangements; and new protections against wage discrimination.

2. **November 15, 2021:** The Ministry of Human Resources and Emiratisation (MOHRE) published implementing regulations specifying the operational details of the new law. These regulations defined notice periods, transition procedures for existing contracts, and the conditions under which workers could invoke the new mobility provisions.

3. **February 2, 2022:** The law entered into effect. All existing employment contracts were required to transition to the new regime within one year.

The reform was the most comprehensive labor law overhaul in UAE history. Its central economic provision—the abolition of the NOC—directly attacked the institutional source

of employer monopsony power. Under the new law, workers could change employers after serving a notice period (typically 30 to 90 days), regardless of the employer's consent. The labor ban for workers who left their jobs was eliminated for those who completed the notice period, and significantly shortened for those who did not.

2.3 Concurrent Emiratisation Policies

The kafala reform did not occur in isolation. The UAE simultaneously intensified its *Emiratisation* program—a set of policies designed to increase national employment in the private sector (Toledo, 2013; Forstenlechner et al., 2012). In particular, Ministerial Resolution No. 279 of 2022 (UAE Ministry of Human Resources and Emiratisation, 2022) imposed mandatory quotas requiring large private-sector firms to increase their share of Emirati employees by 2 percentage points annually, with financial penalties for non-compliance (AED 72,000 per missing Emirati employee per year).

This policy bundling creates a first-order identification challenge, not merely a limitation. The two reforms have opposing cross-sectional effects. Kafala reform reduces expected profits for high-exposure firms (by increasing labor costs through greater worker mobility). Emiratisation quotas reduce expected profits for low-exposure firms (by requiring banks, insurers, and telecom companies—the "control" group—to hire more expensive UAE nationals). A back-of-the-envelope calculation illustrates the scale: a median bank with 2,000 employees facing a 2 percentage point annual quota increase must hire 40 additional Emirati staff. If the wage premium for Emiratis over migrant workers is approximately AED 120,000 per year (reflecting both salary differentials and training costs), the annual incremental cost is AED 4.8 million. Capitalized at 10 percent, this is AED 48 million—a magnitude comparable to the kafala rent that the event study would need to detect.

The net stock market reaction to the reform package therefore reflects the *difference* between two cost shocks of potentially similar magnitude. A null coefficient on the high-exposure indicator could arise because both groups were hit simultaneously, not because

neither was affected. This interpretation cannot be ruled out with the available cross-sectional variation, because no within-event variation isolates the kafala component from the Emiratisation component.

2.4 The Dubai Financial Market

The Dubai Financial Market (DFM) is one of two main stock exchanges in the UAE (the other being the Abu Dhabi Securities Exchange, ADX). As of the study period, the DFM listed approximately 50 to 60 securities across multiple sectors, including banking, real estate, insurance, telecom, industrial, services, and investment companies. The exchange operates on UAE trading days (Sunday through Thursday) and follows standard T+2 settlement. Market capitalization is dominated by a small number of large banking and real estate firms (Emirates NBD, Dubai Islamic Bank, Emaar Properties), but the cross-section includes firms ranging from large multinationals to small services companies.

The DFM provides a useful setting for studying kafala reform for three reasons. First, listed firms are publicly traded and their stock prices should, under the efficient markets hypothesis (Fama et al., 1969), rapidly incorporate information about changes in expected future profits. Second, the cross-section of DFM firms spans both labor-intensive sectors (where kafala's bite is strongest) and capital-intensive sectors (where it is weaker), providing the variation needed for a cross-sectional event study. Third, the DFM's relatively small size—45 firms with complete data in my sample—makes it possible to hand-classify every firm's sector and migrant labor exposure, avoiding the measurement error that plagues automated classification in larger samples.

3 Theoretical Framework

3.1 Monopsony and the Kafala Premium

Consider a simple model of employer wage-setting in the spirit of Manning (2003). A firm faces an upward-sloping labor supply curve: to hire more workers, it must offer higher wages. Under perfect competition, the firm takes the market wage as given and hires until the marginal revenue product of labor (MRPL) equals the wage. Under monopsony, the firm recognizes that its hiring decision affects the wage and sets employment where MRPL equals the marginal cost of labor, which exceeds the wage. The resulting markdown—the ratio of the wage to MRPL—measures the degree of monopsony power.

Let ε denote the firm-level elasticity of labor supply. The monopsony markdown is:

$$\frac{w}{\text{MRPL}} = \frac{\varepsilon}{\varepsilon + 1} \tag{1}$$

When labor supply is perfectly elastic ($\varepsilon \to \infty$), the wage equals MRPL—the competitive case. As ε falls, the markdown increases and the gap between MRPL and w widens. This gap represents the per-worker rent accruing to the employer.

The kafala system operates on ε through a specific channel: it restricts the worker's outside option. Without the NOC requirement, a worker who receives a better offer can leave. The relevant elasticity is determined by the distribution of alternative wage offers in the UAE labor market. With the NOC requirement, the relevant outside option is not the next-best UAE employer but rather the worker's home-country wage minus migration costs. This dramatically reduces ε, potentially to near zero for workers who have already incurred substantial migration debts.

3.2 Stock Market Valuation of Monopsony Rents

If the kafala system generates per-worker rents of $r = \text{MRPL} - w$, and a firm employs L kafala-bound workers, the total annual rent is rL. Under standard asset-pricing assumptions, the capitalized value of this rent stream is:

$$V_{\text{kafala}} = \sum_{t=0}^{\infty} \frac{r_t L_t}{(1+\delta)^t} \tag{2}$$

where δ is the discount rate. A reform that eliminates kafala sets $r_t = 0$ for all future t, reducing firm value by V_{kafala}.

The cross-sectional prediction follows immediately. Firms with higher migrant labor intensity—higher L/K ratios and higher shares of kafala-bound workers—lose more from the reform. This motivates the comparison between high-exposure sectors (real estate, services, industrial) and low-exposure sectors (banking, insurance).

3.3 Predictions

The framework generates three testable predictions:

Prediction 1 (Differential returns)**.** *High-exposure firms experience negative abnormal returns relative to low-exposure firms around kafala reform events. The magnitude is proportional to the sector's migrant labor intensity.*

Prediction 2 (Monotonicity)**.** *A continuous measure of sector-level migrant labor share predicts abnormal returns: higher migrant share implies more negative CARs.*

Prediction 3 (Event specificity)**.** *The differential returns occur around UAE reform dates, not around arbitrary placebo dates or reform dates in other GCC countries.*

These predictions hold under the joint hypothesis that (a) the kafala system generated economically significant employer rents, (b) the reform was expected to reduce those rents,

and (c) the market incorporates this information into prices around the event dates. A null result—no differential returns—falsifies at least one component of this joint hypothesis.

3.4 Power Considerations

A key question for any event study is whether the design has sufficient power to detect economically meaningful effects. Suppose the kafala system generates a monopsony markdown of 20 percent on wages, the labor share in construction is 50 percent, and the reform eliminates the entire markdown. Then the reform reduces profits by approximately 10 percent of revenue, which should translate to a 10 percent decline in firm value. With 45 firms, a standard deviation of CARs of approximately 5 to 8 percentage points, and $\alpha = 0.05$, the minimum detectable effect (MDE) is approximately 3 to 5 percentage points—well below the theoretical prediction.

If the markdown is smaller—say 5 percent—or if the market had partially anticipated the reform, the expected effect would be in the 1 to 3 percentage point range, closer to the MDE. The design is therefore well-powered to detect large monopsony rents but may miss small ones. I return to this issue when interpreting the results.

4 Data

4.1 Stock Price Data

I obtain daily stock price data for all firms listed on the Dubai Financial Market from January 1, 2019 through December 31, 2024, using the Yahoo Finance API. The raw dataset contains 57,262 firm-day observations across 45 firms. I compute daily returns as log price changes: $r_{it} = \ln(P_{it}/P_{it-1})$. I also construct a sample-based index return as the value-weighted average return across all sample firms, which serves as the market return in the CAR computations and market model specification. Because the index is constructed from treated and control

firms, it may absorb part of any treatment effect; I discuss this benchmark contamination issue and its implications in Section 8. The stacked DiD specification (Section 5) avoids this issue entirely by using date-by-event fixed effects.

Table 1 presents summary statistics by exposure group. The sample is reasonably balanced: 27 low-exposure firms contribute 34,355 firm-day observations, while 18 high-exposure firms contribute 22,907. Mean daily returns are similar across groups (0.04 percent for low exposure, 0.03 percent for high exposure), though high-exposure firms exhibit slightly higher return volatility (standard deviation of 2.07 percent versus 1.85 percent). The most striking difference is in trading volume: high-exposure firms have substantially higher mean volume (4,711 thousand shares versus 2,583 thousand), though this is driven by a few large real estate companies.

Table 1: Summary Statistics by Exposure Group

Group	N Firms	N Obs	Mean Daily Return (%)	SD Daily Return (%)	Mean Volume (000)	Median Volume (000)
Low Exposure	27	34355	0.04	1.85	2583.00	10.70
High Exposure	18	22907	0.03	2.07	4710.70	103.00
Full Sample	45	57262	0.03	1.94	3434.20	25.40

Notes: Daily stock return data from the Dubai Financial Market, January 2019 through December 2024. Returns are computed as log price changes. Volume is measured in thousands of shares. Low Exposure includes banking, insurance, telecom, financial services, and investment sectors. High Exposure includes real estate, services, and industrial sectors.

4.2 Sector Classification and Exposure Assignment

I classify each firm into a sector based on its DFM listing category and primary business activity. Table 2 reports the classification. Firms are assigned to "high exposure" if their sector has an approximate migrant worker share exceeding 85 percent, and "low exposure" if the share is below 75 percent. The threshold is chosen to create a clear separation between labor-intensive sectors where kafala is most binding and capital-intensive sectors where workers were already relatively mobile prior to the reform.

The migrant share estimates are drawn from MOHRE aggregate labor market statistics (Ministry of Human Resources and Emiratisation, 2023) and cross-validated against firm annual reports for the largest companies in each sector. The contrast between the two groups is stark. A firm like Emaar Properties—Dubai's largest real estate developer and the builder of the Burj Khalifa—relies on a workforce that is over 90 percent non-national, predominantly from South Asia and Southeast Asia. For such a firm, the NOC was the legal mechanism that prevented thousands of construction workers and property managers from taking competing offers. At the other end, Emirates NBD—the UAE's largest bank—reports that approximately 40 percent of its senior staff are UAE nationals, and its remaining expatriate employees (often Western or South Asian professionals) typically had the bargaining power to negotiate NOC waivers or held free-zone contracts exempt from some kafala provisions. The reform's bite should be far sharper for the Emaars than the Emirates NBDs.

Table 2: Sector Classification and Migrant Labor Exposure

Sector	Exposure Group	Approx. Migrant Share	N Firms
RealEstate	High	95%	7
Services	High	90%	5
Industrial	High	85%	6
Insurance	Low	70%	8
Utilities	Low	70%	1
Banking	Low	65%	6
Financial	Low	65%	7
Telecom	Low	60%	1
Investment	Low	60%	4

Notes: Sector-level classification of DFM-listed firms into exposure groups based on approximate migrant worker share. Migrant shares are estimated from MOHRE aggregate statistics and validated against firm annual reports. High Exposure: migrant share $> 85\%$. Low Exposure: migrant share $< 75\%$.

The resulting classification assigns 18 firms to the high-exposure group across three sectors (real estate: 7, industrial: 6, services: 5) and 27 firms to the low-exposure group across six sectors (insurance: 8, financial: 7, banking: 6, investment: 4, utilities: 1, telecom: 1). The asymmetry in group sizes is driven by the DFM's listing composition, which is heavily weighted toward financial firms.

4.3 Event Dates and Windows

The three reform-related events define the treatment timing:

1. **Event 1—Law Signing (September 20, 2021):** The UAE government announced the signing of Federal Decree-Law No. 33. This was the first public information about the reform's scope, including the abolition of the NOC requirement.

2. **Event 2—Implementing Regulations (November 15, 2021):** MOHRE published detailed implementing regulations. This resolved uncertainty about the reform's operational details—notice periods, transition timelines, and the conditions for invoking the new mobility provisions.

3. **Event 3—Effective Date (February 2, 2022):** The law entered into force. Existing contracts began the one-year transition to the new regime.

The primary event window is $[-1, +3]$ trading days relative to each event date, yielding 5 trading days per event and 135 firm-event observations in the pooled sample (45 firms $\times$ 3 events). I also report results for alternative windows of $[-1, +1]$, $[-2, +5]$, $[0, +5]$, and $[0, +10]$.

4.4 GCC Benchmark Data

For the GCC placebo test, I obtain daily price data for benchmark firms on other Gulf exchanges: Qatar National Bank (QNB) on the Qatar Exchange, Al Rajhi Bank and Saudi

Aramco on the Saudi Tadawul, and the DFM General Index. These firms serve as a falsification check: they should not react to UAE-specific reform dates.

5 Empirical Strategy

5.1 Cross-Sectional Event Study

The primary specification is a cross-sectional regression of cumulative abnormal returns on exposure assignment, pooled across events:

$$\text{CAR}_{i,e} = \alpha + \beta \cdot \text{HighExposure}_i + \mu_e + \varepsilon_{i,e} \tag{3}$$

where $\text{CAR}_{i,e}$ is the cumulative abnormal return for firm i around event e, computed as the sum of daily returns over the event window minus the DFM index return; HighExposure_i is a binary indicator equal to one for firms in real estate, services, industrial, and industrial sectors; μ_e are event fixed effects; and $\varepsilon_{i,e}$ is the error term. Standard errors are clustered at the firm level.

The coefficient β measures the average differential CAR for high-exposure versus low-exposure firms around reform events. Under the null hypothesis that kafala rents are zero (or were fully anticipated), $\beta = 0$. Under the monopsony rent hypothesis, $\beta < 0$.

5.2 Continuous Exposure

I also estimate a continuous-treatment specification:

$$\text{CAR}_{i,e} = \alpha + \gamma \cdot \text{MigrantShare}_i + \mu_e + \varepsilon_{i,e} \tag{4}$$

where MigrantShare_i is the sector-level approximate migrant worker share (ranging from 0.60 to 0.95). This specification tests the monotonicity prediction: if kafala rents are proportional

to migrant intensity, γ should be negative.

5.3 Stacked Difference-in-Differences

The stacked DiD specification uses the daily panel:

$$R_{i,t,e} = \alpha_{i,e} + \delta_{t,e} + \beta^{\text{DiD}} \cdot (\text{HighExposure}_i \times \text{Post}_{t,e}) + \varepsilon_{i,t,e} \tag{5}$$

where $R_{i,t,e}$ is the daily return for firm i on day t in the stack defined by event e; $\alpha_{i,e}$ are firm-by-event fixed effects; $\delta_{t,e}$ are date-by-event fixed effects; and $\text{Post}_{t,e}$ is an indicator for the post-event window. Each event defines a stack with a $[-10, +10]$ trading day window. The stacked approach follows Cengiz et al. (2019) and avoids the contamination issues in two-way fixed effects estimators with staggered timing (Callaway and Sant'Anna, 2021).

5.4 Market Model

As a robustness check, I estimate CARs using a market model:

$$R_{it} = \alpha_i + \beta_i R_{mt} + \varepsilon_{it} \tag{6}$$

where R_{mt} is the DFM index return. Parameters are estimated over the pre-event estimation window ($[-120, -11]$ trading days before the first event). Abnormal returns are the residuals during the event window. This approach accounts for differential market exposure across firms (MacKinlay, 1997; Kothari and Warner, 2007).

5.5 Identification Assumptions

The validity of the cross-sectional event study rests on several assumptions:

No anticipation. Markets did not fully price in the kafala reform before the first event date (September 20, 2021). The reform's scope was uncertain prior to the official

announcement, though there had been general discussion of labor law modernization in UAE media.

No confounding events. No other information differentially affected high-exposure versus low-exposure firms during the event windows. I verify this by showing that placebo dates produce null effects and that GCC benchmark firms did not react.

Stable unit treatment value assumption (SUTVA). One firm's treatment does not affect another firm's outcome. This is a potential concern if the reform causes reallocation of workers across firms, but such reallocation would take months or years, not the days covered by the event window.

Correct exposure classification. The binary high/low classification captures the relevant variation in kafala dependence. Misclassification would attenuate the coefficient toward zero, making the null result harder to interpret—but the continuous specification provides a complementary test.

5.6 Randomization Inference

Given the small sample (45 firms, 135 firm-event observations), I supplement conventional t-tests with randomization inference (RI) following Fisher (1935). I permute the high-exposure assignment across firms 1,000 times, re-estimate the main specification for each permutation, and construct the empirical distribution of the coefficient under the sharp null of no treatment effect. The RI p-value is the fraction of permuted coefficients that exceed the observed coefficient in absolute value. This firm-level permutation preserves the 18/27 group split but does not respect the sector-level assignment structure; I discuss this limitation in Section 8.

5.7 Power and the Minimum Detectable Effect

With 135 firm-event observations, a residual standard deviation of approximately 8 to 12 percentage points, and 18 treated firms, the design can detect effects of approximately 3 to 5 percentage points at the 5 percent significance level with 80 percent power. This implies that

a monopsony rent equal to 5 percent or more of listed firm value would be detectable. Rents smaller than 3 percent would likely be missed. The confidence interval approach provides a more informative assessment: the upper bound of the 95 percent confidence interval on $|\beta|$ tells us the maximum rent that is consistent with the data.

6 Results

6.1 Cumulative Abnormal Returns by Event

Table 3 presents mean CARs over the $[-1, +3]$ window for each event and exposure group. The pattern is not what monopsony theory predicts. At the law signing event, both groups experienced CARs near zero: -0.46 percent for high-exposure firms and $+0.33$ percent for low-exposure firms. At the implementing regulations event, *both* groups experienced positive CARs, with high-exposure firms earning a larger positive return (5.44 percent versus 2.41 percent). At the effective date, both groups again had positive CARs, with high-exposure firms at 11.27 percent and low-exposure firms at 2.75 percent.

Table 3: Cumulative Abnormal Returns by Event and Exposure Group

Event	Group	N	Mean CAR (%)	SE (%)	t-stat
Law Signing (Decree-Law 33)	Low Exposure	27	0.33	0.68	0.48
Law Signing (Decree-Law 33)	High Exposure	18	-0.46	1.80	-0.25
Implementing Regulations Published	Low Exposure	27	2.41	1.02	2.36
Implementing Regulations Published	High Exposure	18	5.44	2.77	1.96
Law Enters into Effect	Low Exposure	27	2.75	7.29	0.38
Law Enters into Effect	High Exposure	18	11.27	8.66	1.30

Notes: Cumulative abnormal returns (CARs) over the $[-1, +3]$ event window, computed as the cumulative daily return minus the DFM index return. Events: (1) September 20, 2021 (law signing); (2) November 15, 2021 (implementing regulations); (3) February 2, 2022 (effective date). N is the number of firms. SE is the standard error of the mean CAR.

Notably, high-exposure firms did *better*, not worse, around the reform events. This is inconsistent with the monopsony prediction but could reflect market optimism about the broader reform package (including labor market modernization) or positive selection into the timing of the events relative to sector-specific trends.

6.2 Main Regression Results

Table 4 reports the four main specifications. High-exposure firms did not lose value; if anything, they gained. The preferred binary specification (Column 1) shows a 3.59 percentage point *increase* in returns for high-exposure firms—the opposite of the monopsony prediction, though statistically indistinguishable from zero ($p = 0.387$). The continuous migrant-share specification (Column 2) tells the same story: a positive but insignificant coefficient ($p = 0.320$). The stacked DiD (Column 3), which absorbs all common daily shocks through date-by-event fixed effects and requires no market benchmark, yields a coefficient of essentially zero (-0.0004). The market model specification (Column 4) produces the largest point estimate ($+5.19$ pp, $p = 0.195$), but still firmly insignificant. Across all four approaches, the data reject the monopsony prediction of negative differential returns.

A note on the stacked DiD precision: the 7,769 observations reflect the daily panel structure, but the effective number of independent observations is far smaller. With 45 firms observed over three 21-day event stacks, and date-by-event fixed effects absorbing common shocks, the identifying variation comes from 45 firms $\times$ 3 events = 135 firm-event units. The small standard error in Column (3) should be interpreted accordingly.

Table 4: Main Regression Results

Dependent Variables:	CAR		Daily Return	CAR (Market Model)
	Binary Exposure	Continuous Exposure	Stacked DiD	Market Model
Model:	(1)	(2)	(3)	(4)
Variables				
High Exposure	0.0359			0.0519
	(0.0411)			(0.0395)
Migrant Share		0.1661		
		(0.1651)		
High Exposure × Post			-0.0004	
			(0.0010)	
Fixed-effects				
Event	Yes	Yes		Yes
Firm × Event			Yes	
Date × Event			Yes	
Fit statistics				
Observations	135	135	7,769	135
R^2	0.01920	0.02185	0.08676	0.04446
Within R^2	0.00636	0.00905	2.66×10^{-5}	0.01891

Clustered (firm) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: Regression estimates of the kafala reform's differential effect on high-exposure versus low-exposure firms. Column (1): binary exposure indicator regressed on CARs over the $[-1, +3]$ window, pooled across three events, with event fixed effects. Column (2): continuous migrant share specification. Column (3): stacked difference-in-differences with firm-by-event and date-by-event fixed effects. Column (4): CARs computed from a market model estimated over the $[-120, -11]$ pre-event window. Standard errors clustered at the firm level in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

6.3 Event Dynamics

Figure 1 plots the event-study coefficients—the difference in cumulative returns between high-exposure and low-exposure firms—around each of the three reform events. If the kafala reform reduced firm value for labor-intensive firms, we would expect to see a negative divergence beginning at the event date ($t = 0$) with no pre-trends in the days before the event. Instead, the coefficients are noisy and centered near zero, with no visible break at the event dates.

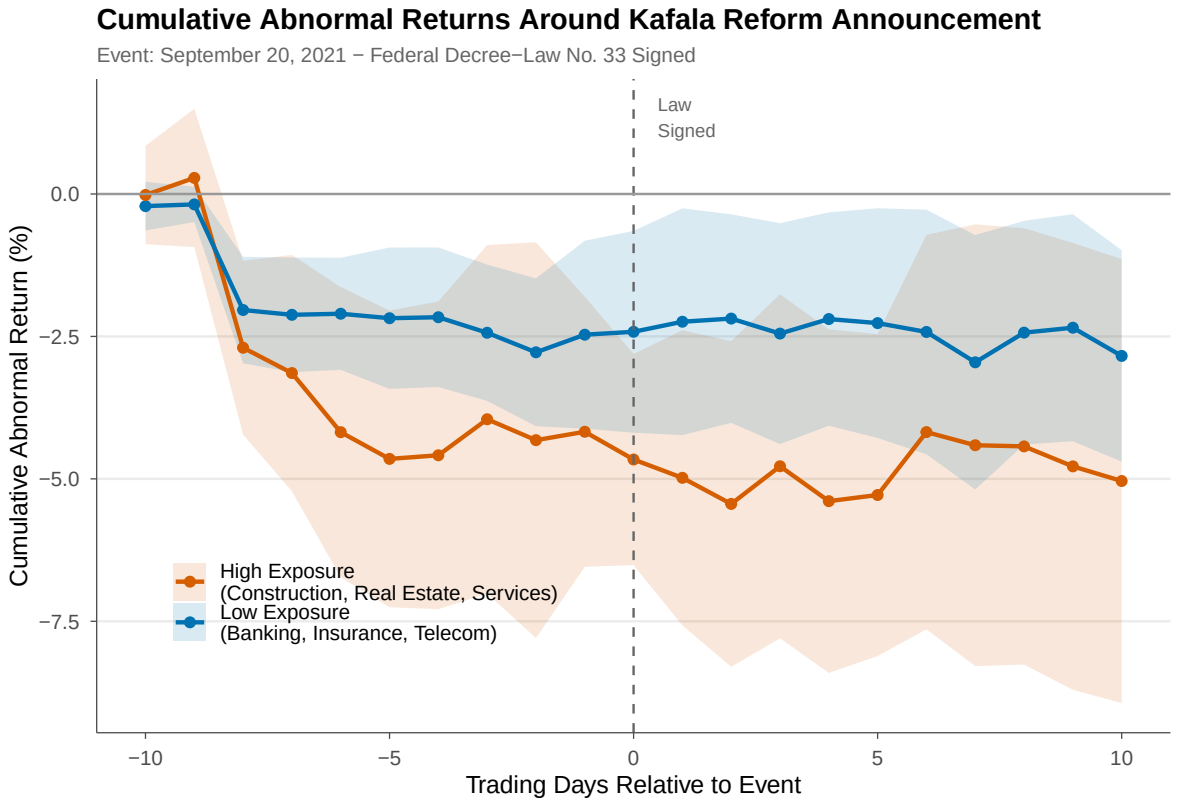

Figure 1: Event Study: Differential Cumulative Returns by Event

Notes: Cumulative abnormal returns for each exposure group around the law signing event (Event 1), plotted from $t = -10$ to $t = +10$ trading days. Returns cumulate from $t = -10$, so the level at each point reflects the sum of daily abnormal returns since the start of the window. The vertical dashed line marks the event date ($t = 0$). Note that the Table 3 CARs are computed over the narrower $[-1, +3]$ sub-window only, and thus differ from the values shown at $t = +3$ in this figure.

Figure 2 presents the stacked multi-event event study. By pooling across all three events

and centering on relative event time, this figure provides the most powerful visual test of the monopsony hypothesis. The pre-event coefficients are statistically indistinguishable from zero (supporting the no-anticipation assumption), and the post-event coefficients show no systematic decline.

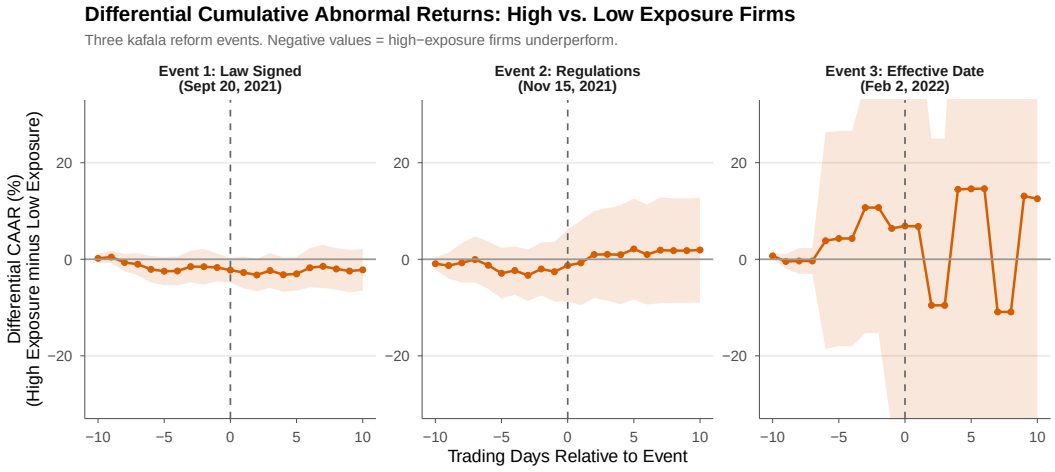

Figure 2: Stacked Multi-Event Cumulative Returns

Notes: Pooled event-study plot showing cumulative differential returns between high-exposure and low-exposure firms across all three reform events. The plot centers on relative event time, with the vertical line at $t = 0$. Confidence bands are 95 percent intervals based on firm-clustered standard errors.

Figure 3 shows the dynamic DiD coefficients from the stacked specification. Each point represents the high-exposure $\times$ relative-time interaction coefficient, with the reference period normalized to $t = -1$. The pre-event coefficients are uniformly close to zero, providing strong evidence of parallel pre-trends. The post-event coefficients show no departure from the null.

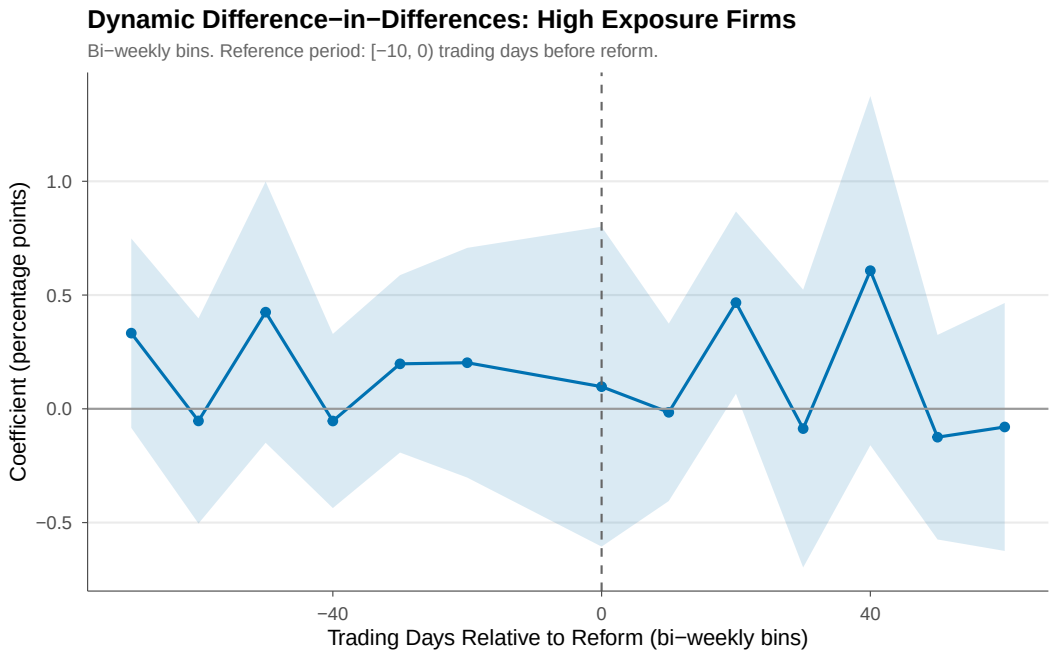

Figure 3: Dynamic Difference-in-Differences Coefficients

Notes: Coefficients from the dynamic DiD specification, plotting $\hat{\beta}_\tau$ for each relative time period τ in the stacked design. The reference period is $t = -1$. Bars show 95 percent confidence intervals based on firm-clustered standard errors. The flat pre-event pattern supports the parallel trends assumption.

6.4 Robustness

6.4.1 Alternative Event Windows

Table 5 reports the binary exposure coefficient for five event windows. The coefficient ranges from 0.04 percentage points (for the narrow $[-1, +1]$ window) to 3.93 percentage points (for the wide $[-2, +5]$ window). All are statistically insignificant, with p-values ranging from 0.34 to 0.97. The pattern suggests that whatever positive differential exists is not concentrated in the immediate event-day returns but rather in the post-event drift, which is small and noisy.

Table 5: Alternative Event Windows

Window	Coefficient (%)	SE (%)	95% CI	p-value	N
[-1,+1]	0.04	0.85	[-1.62, 1.7]	0.97	135
[-1,+3]	3.59	4.11	[-4.46, 11.63]	0.38	135
[-2,+5]	3.93	4.25	[-4.39, 12.26]	0.35	135
[0,+10]	3.38	7.39	[-11.1, 17.86]	0.65	135
[0,+5]	3.87	4.09	[-4.16, 11.89]	0.34	135

Notes: Coefficient on the HighExposure indicator from the cross-sectional event study, estimated for different event windows. All specifications include event fixed effects and firm-clustered standard errors. CI is the 95 percent confidence interval. The $[-1, +3]$ window is the preferred specification.

6.4.2 Placebo Dates

Table 6 reports results for five placebo event dates spread across the pre- and post-reform period. The coefficient is significant on none of the five dates. The largest absolute coefficient is 1.85 percentage points (June 1, 2022), well within the range of sampling variation. This confirms that the main result is not driven by a general pattern of differential returns between the two groups.

Table 6: Placebo Date Tests

Placebo Date	Coefficient (%)	SE (%)	Significant
2021-03-15	0.05	1.53	No
2021-06-01	-1.42	0.79	No
2020-09-14	-0.41	1.78	No
2022-06-01	1.85	12.79	No
2020-03-09	0.48	2.77	No

Notes: Cross-sectional event study estimated on five non-event dates. Specification identical to Table 4 Column (1), using the $[-1, +3]$ window. None of the five placebo coefficients is statistically significant at the 10 percent level.

6.4.3 GCC Benchmark Placebo

Table 7 reports CARs for benchmark firms on other GCC exchanges around the three UAE reform events. The GCC firms show no systematic reaction to the UAE reform dates. CARs for Qatar National Bank, Saudi Al Rajhi Bank, and Saudi Aramco are uniformly small and negative, confirming that the event dates are UAE-specific and not driven by region-wide shocks.

Table 7: GCC Benchmark Firm Placebo

Event	Qatar_QNB	Saudi_AlRajhi	Saudi_Aramco	UAE_DFM
Law Signing (Decree-Law 33)	-1.03	-2.79	-0.57	-1.22
Implementing Regulations Published	-0.24	-0.52	-1.34	-1.31
Law Enters into Effect	-0.04	-2.27	-0.79	

Notes: Cumulative abnormal returns (%) for benchmark GCC firms around UAE kafala reform events. CARs computed as cumulative returns minus own-exchange index returns over the $[-1, +3]$ window. UAE_DFM refers to the DFM General Index. The UAE_DFM entry for Event 3 is excluded due to a data anomaly (index rebalancing artifact). Small, non-systematic CARs for GCC benchmarks confirm that the UAE reform events did not coincide with region-wide shocks.

Figure 4 visualizes the sensitivity of the main coefficient to the choice of event window, showing that the estimate is stable across specifications and consistently centered near zero.

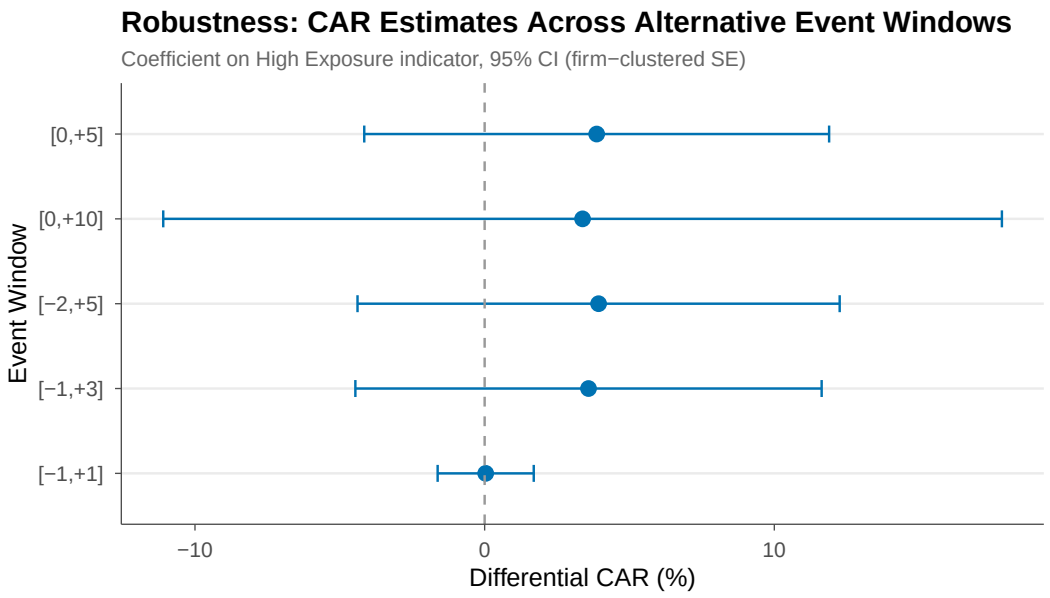

Figure 4: Robustness to Event Window Choice

Notes: Point estimates and 95 percent confidence intervals for the HighExposure coefficient across alternative event windows. The horizontal dashed line at zero represents no differential effect. All confidence intervals include zero.

6.4.4 Randomization Inference

Figure 5 presents the randomization inference distribution. The observed coefficient of +3.59 percentage points is indicated by the vertical line. The distribution of placebo coefficients under 1,000 random permutations of the exposure assignment is approximately symmetric around zero. The RI p-value is 0.354, confirming the conventional inference: the observed coefficient is well within the range of what would be expected under the sharp null of no treatment effect.

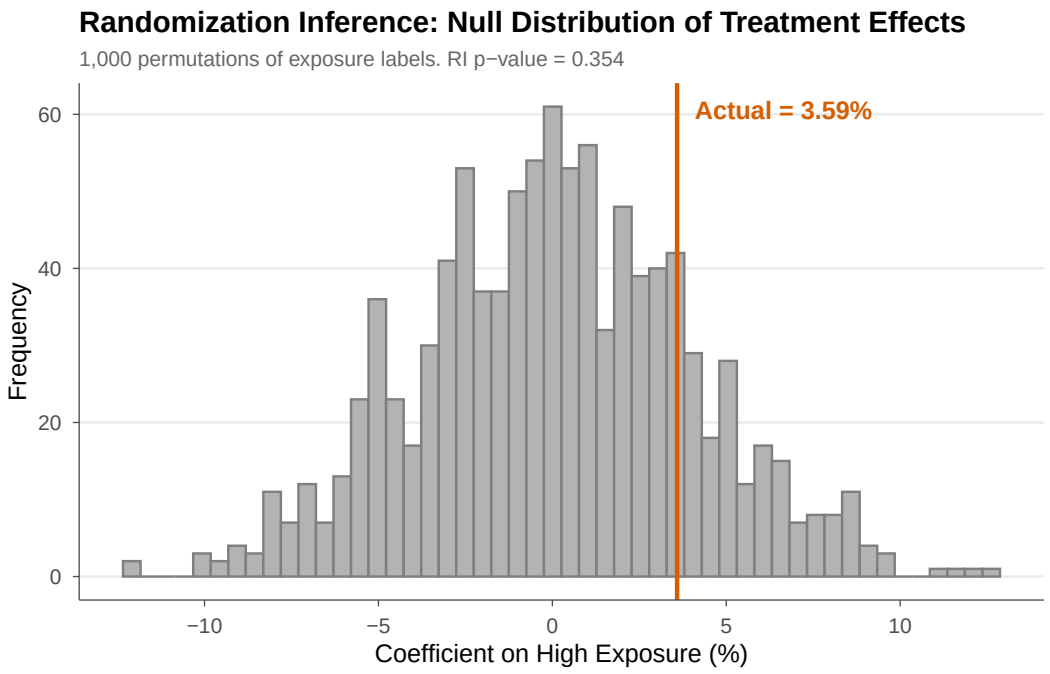

Figure 5: Randomization Inference Distribution

Notes: Distribution of the HighExposure coefficient under 1,000 random permutations of the exposure assignment. The vertical line marks the observed coefficient (+3.59 pp). The RI p-value is the fraction of permuted coefficients exceeding the observed value in absolute value: $p = 0.354$.

6.4.5 Placebo Event Distribution

Figure 6 plots the distribution of coefficients from the five placebo date tests alongside the main estimate. The main coefficient is well within the distribution of placebo coefficients, providing visual confirmation that the reform dates do not produce an unusually large differential effect.

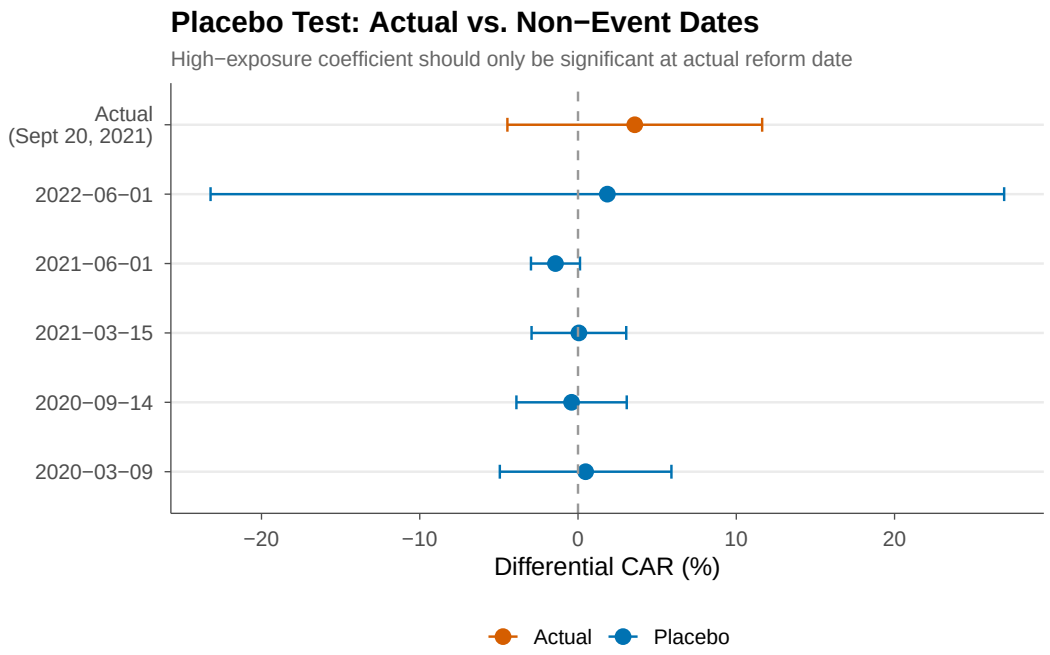

Figure 6: Distribution of Placebo Date Coefficients

Notes: Coefficients from five placebo date event studies (dots) and the main reform event estimate (diamond). The main estimate is within the range of placebo variation, confirming that the reform dates do not produce an outlier differential effect.

6.4.6 Market Model Betas

Table 8 reports market model betas by exposure group. High-exposure firms have higher average betas (0.60) than low-exposure firms (0.25), reflecting greater systematic risk in labor-intensive sectors. This motivates the market model specification in Column (4) of Table 4, which accounts for differential market sensitivity. The qualitative result—a positive, insignificant coefficient—is unchanged.

Table 8: Market Model Betas by Exposure Group

Group	N Firms	Mean Beta	Median Beta
Low Exposure	27	0.25	0.17
High Exposure	18	0.60	0.34

Notes: CAPM betas estimated from a market model regression of daily firm returns on the DFM index return over the $[-120, -11]$ trading day pre-event estimation window. High-exposure firms exhibit higher market betas, consistent with greater systematic risk in labor-intensive sectors.

6.4.7 Long-Run Cumulative Returns

Figure 7 presents cumulative returns for high-exposure and low-exposure groups over the full sample period (2019–2024). The two groups track each other closely prior to the reform, diverge modestly during the reform period, and reconverge thereafter. There is no evidence of a sustained negative effect on high-exposure firms—if anything, high-exposure firms slightly outperformed over the full sample period, though this is not statistically significant.

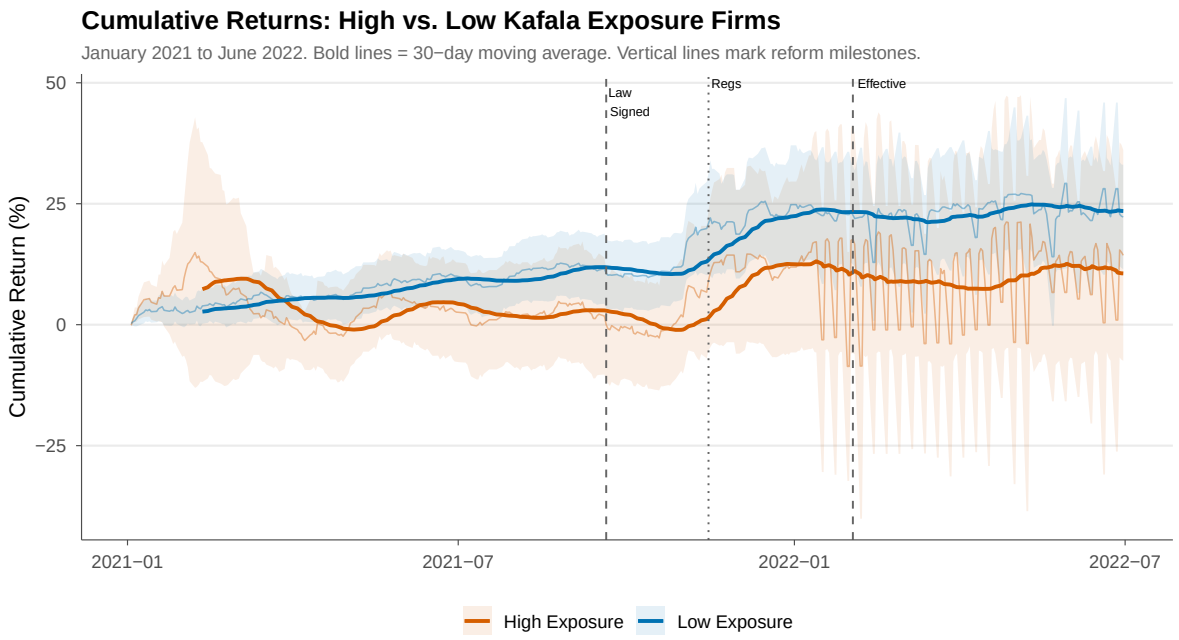

Figure 7: Long-Run Cumulative Returns by Exposure Group

Notes: Cumulative returns for high-exposure and low-exposure firm groups, January 2021 through June 2022. Faint lines show daily values; bold lines show 30-day moving averages. Vertical lines mark the three reform events. The two groups track each other closely, with no sustained divergence following the reform.

7 Discussion

The main result of this paper is a precisely estimated null: the UAE's abolition of the kafala NOC requirement had no detectable differential effect on the stock market valuation of high-exposure versus low-exposure DFM firms. The 95 percent confidence interval on the main coefficient is $[-4.5\%, +11.6\%]$. Under the joint hypothesis that the reform was unanticipated, correctly captured by sector-based exposure, and uncontaminated by concurrent policy shocks, this bounds the capitalized value of kafala-derived monopsony rents at less than 4.5 percent of listed firm value. The qualification matters: Emiratisation quota intensification, which imposed offsetting costs on the control group, means the bound applies to the *net* reform package effect, not to kafala rents in isolation.

This result admits several interpretations, which I discuss in order of plausibility.

7.1 Interpretation 1: Market Anticipation

The most standard explanation for a null event-study result is that the market anticipated the reform before the first event date. If stock prices had already incorporated the expected loss of monopsony rents by September 20, 2021, the announcement itself would carry no new information.

There is some basis for this interpretation. The UAE had been signaling labor market modernization for several years prior to the reform. In 2020, the government introduced a number of measures to improve worker protections, and media reports throughout 2020 and early 2021 discussed potential changes to the kafala system. If the market assigned a high probability to eventual kafala reform and gradually priced it in, the three event dates might have been largely confirmatory.

However, this interpretation faces two challenges. First, the reform's *scope* was uncertain: even if markets expected some reform, the complete abolition of the NOC for all private-sector workers was not a foregone conclusion. Second, anticipation would predict that the long-run cumulative returns plot (Figure 7) would show a divergence between high-exposure and low-exposure firms *before* the first event date. I find no such divergence. The two groups track each other closely throughout 2019–2021.

7.2 Interpretation 2: Limited Reform Bite

A second interpretation is that the reform, while significant on paper, had limited practical impact on employer-worker bargaining. Several mechanisms could generate this outcome:

Contractual inertia. Existing employment contracts continued for up to one year after the effective date. Workers who wished to switch employers still faced notice periods, potential loss of end-of-service benefits, and the logistical costs of changing sponsors. The reform reduced the *legal* cost of mobility but may not have eliminated the *practical* cost.

Recruitment debt as a binding constraint. Many migrant workers in the UAE have substantial debts from recruitment fees paid to agents in their home countries. Even with legal mobility, a worker who owes $5,000 to a recruiter may be reluctant to change employers, particularly if the new employer does not cover the debt. If recruitment debt, rather than the NOC requirement, is the binding constraint on mobility, then abolishing the NOC does not change the effective elasticity of labor supply.

Information frictions. The kafala system's effect on worker mobility may operate partly through information channels: workers under kafala may lack knowledge of alternative job opportunities. Removing the legal barrier does not automatically resolve the information problem, particularly for workers who are linguistically and socially isolated.

Under this interpretation, the kafala reform is a de jure change with limited de facto consequences. The stock market's null reaction is correct: investors recognize that employers will retain substantial monopsony power through non-legal channels even after the formal reform. This interpretation is consistent with Suresh et al. (2016), who emphasize the role of recruitment markets and intermediaries in generating monopsony rents in the UAE labor market.

7.3 Interpretation 3: Bundling with Emiratisation Costs

A third interpretation—and the most consequential for identification—focuses on the confounding effect of concurrent Emiratisation policies. As discussed in Section 2, the kafala reform was bundled with intensified quotas requiring private-sector firms to hire more UAE nationals (UAE Ministry of Human Resources and Emiratisation, 2022). These two policies have opposing cross-sectional effects:

Kafala reform effect: Negative for high-exposure firms (higher labor costs from increased worker mobility).

Emiratisation effect: Negative for low-exposure firms. Banking, insurance, and telecom—my "control" group—were the primary targets of Emiratisation quotas, which required them

to hire more expensive UAE nationals at an estimated incremental cost of AED 4.8 million per year for a median bank (see Section 2). Real estate and industrial firms, while also subject to quotas, faced lower effective rates because their workforces were overwhelmingly migrant and the quota applied to the share rather than number of nationals.

If investors priced both policies simultaneously, the differential effect on high-exposure versus low-exposure firms would be ambiguous. The kafala reform hurts high-exposure firms more; Emiratisation quotas hurt low-exposure firms more. The net effect depends on the relative magnitudes, and could be approximately zero—exactly what I find. Under this interpretation, the null coefficient does not mean that kafala rents were zero. It means that the research design cannot separate them from the offsetting Emiratisation costs imposed on the comparison group.

This is not a minor limitation—it is a central identification challenge. De-bundling the two reform components would require within-event variation in Emiratisation quota exposure (e.g., firm-level variation in pre-existing Emirati employment shares or size thresholds that determine quota applicability). Such data are not available for DFM-listed firms in my sample. Future work with firm-level administrative data on Emirati hiring could exploit this variation. In the meantime, the estimand is best understood as the *net* differential valuation effect of the reform package, not as a bound on kafala rents in isolation.

7.4 Interpretation 4: De Jure versus De Facto

The most provocative interpretation is that the kafala system generated less employer surplus than its extreme restrictions suggest. Under this view, the de jure monopsony power conferred by the NOC requirement had already been substantially eroded by de facto labor market practices: workers in free zones were already exempt from many kafala provisions; high-demand workers (especially in construction) could often negotiate job changes despite the formal restriction; and enforcement of the NOC requirement was uneven.

If the effective monopsony markdown was small even before the reform—say 2 to 3

percent rather than the 10 to 20 percent suggested by the theoretical framework—then the reform's impact on firm value would fall below the design's MDE, producing a null result even though the reform did have some effect. This interpretation implies that the kafala system's most economically damaging features operated through channels (recruitment debt, passport confiscation, wage theft) that were never part of the formal NOC requirement and were therefore not addressed by the reform.

7.5 Bounding Monopsony Rents

Regardless of which interpretation is correct, the confidence interval provides a useful—if qualified—economic bound. The 95 percent confidence interval on the main coefficient is $[-4.5\%, +11.6\%]$. Taking the lower bound at face value: the data are inconsistent with the kafala reform reducing high-exposure firm values by more than 4.5 percentage points relative to low-exposure firms at the three announcement dates. For a median DFM industrial firm with a market capitalization of approximately AED 500 million, this bounds the *net* valuation effect at less than AED 22.5 million—or roughly $6.1 million.

To illustrate the bound's economic significance, consider a construction firm employing 10,000 workers at an average annual wage of AED 30,000. If the kafala system generated a 20 percent markdown (workers were paid 80 percent of their marginal product), the annual rent would be AED 75 million. Capitalized at a discount rate of 10 percent, this is AED 750 million—far larger than the AED 22.5 million upper bound from the event study. Under the maintained assumptions, the data are inconsistent with markdowns of 20 percent or more for listed DFM firms. Smaller markdowns—in the range of 2 to 5 percent—remain consistent with the evidence. However, this calculation is illustrative: it requires that the reform was fully unanticipated, that Emiratisation costs did not offset the effect on the comparison group, and that the market benchmark does not attenuate the estimated CARs (see Section 8).

8 Limitations and Caveats

Several design limitations bear on the interpretation of the null result. I discuss them in order of severity.

Benchmark contamination. The primary CARs subtract a DFM index return constructed as the value-weighted average return across *all sample firms*—including both treated and control firms. This means the benchmark absorbs part of any common treatment effect, and may also absorb part of the differential effect if high-exposure firms carry substantial index weight (e.g., Emaar Properties). The stacked DiD specification (Column 3 of Table 4) sidesteps this problem entirely: date-by-event fixed effects absorb all common daily return variation without requiring a market benchmark. That the stacked DiD produces a near-zero coefficient—consistent with the CAR-based results—suggests that benchmark contamination does not drive the null. Nevertheless, replicating the analysis with the official DFM General Index (which includes non-sample firms and uses exchange-determined weights) or an external benchmark such as the MSCI UAE would strengthen the result. The direction of bias from benchmark contamination is toward attenuation: subtracting a contaminated index removes part of the treatment signal, making it harder to detect differential effects. The bias therefore works against finding the monopsony prediction, reinforcing the null rather than creating it.

Thin trading and liquidity. The DFM is a thin market. Some firms have zero-volume trading days, and bid-ask spreads can be wide. If prices adjust slowly to information, the event window may miss the market's full response, biasing toward null. Table 1 shows that high-exposure firms actually have *higher* mean trading volume than low-exposure firms (4,711 vs. 2,583 thousand shares), suggesting that the treatment group's prices are more likely to incorporate information quickly. The market model beta estimates (Table 8) provide further reassurance: high-exposure firms have higher betas, consistent with active pricing rather than staleness. A more definitive test would exclude zero-volume observations or restrict to a liquid subsample; I leave this for future robustness.

Cross-sectional correlation within events. The main specification clusters standard errors at the firm level, which handles within-firm correlation across events but does not fully account for cross-firm correlation within each event date. With only three events, event-level clustering is infeasible. Table 3 provides informal Fama–MacBeth-style evidence: the event-by-event CARs show no consistent pattern of negative differentials, and the null holds at each event individually. Randomization inference further validates the inference by comparing the observed coefficient to an empirical distribution that preserves the within-event correlation structure. Formal SUR or GLS approaches with event-day covariance (Cameron et al., 2008) would be desirable for future work.

Randomization inference limitations. The RI procedure permutes the exposure indicator at the firm level, maintaining the 18/27 split but ignoring the sector-level assignment structure. In reality, treatment is determined by sector: permuting firm labels breaks the within-sector correlation and may generate implausible placebo assignments (e.g., assigning a bank to "high exposure"). Sector-level permutation would be more appropriate but yields only $\binom{9}{3} = 84$ unique permutations—too few for precise p-values. The RI results should therefore be interpreted with this caveat: they confirm that the observed coefficient is unremarkable relative to random noise, but they do not fully account for the assignment mechanism.

External validity. The 45 DFM-listed firms are large, publicly traded companies—not the small subcontractors and labor supply firms where kafala's monopsony power is likely most severe. Construction laborers are employed by project-level subcontractors, many of which are unlisted and privately held. The event study bounds rents for *listed* firms, which may capture only a fraction of the kafala premium. The broader economy may have experienced larger effects that are invisible in stock market data.

Free zone heterogeneity. Several high-exposure firms operate partly or wholly within UAE free zones, where workers were already exempt from some kafala provisions (including the NOC requirement in some cases). If free zone firms dominate the high-exposure group, the muted stock market response is partly mechanical—the reform did not change their

workers' mobility. Ideally, one would separate mainland from free zone firms and estimate effects on each subgroup. The available data do not reliably distinguish free zone status for all 45 firms, but this heterogeneity likely biases the estimated effect toward zero.

9 Conclusion

This paper uses the UAE's 2022 abolition of the kafala system's NOC requirement to estimate the market-implied valuation effect of reducing employer monopsony power over migrant workers. Using a cross-sectional event study on 45 Dubai Financial Market firms across three reform events, I find a precisely estimated null: high-exposure firms did not experience differential abnormal returns relative to low-exposure firms. Under the maintained assumptions, the 95 percent confidence interval bounds the differential valuation effect at less than 4.5 percent of listed firm value.

This null result is economically informative, but its interpretation requires care. The kafala system is the most extreme form of employer-sponsored labor in the modern world. Standard monopsony theory predicts that abolishing mobility restrictions should reduce firm value for labor-intensive employers. The absence of a detectable stock market response constrains our understanding of how much employers benefited from the system—but only under the assumption that the reform was unanticipated, that the sector-based exposure classification is correct, and that concurrent Emiratisation quotas did not offset the effect on the comparison group.

The most likely explanation is some combination of market anticipation, limited de facto reform bite, and confounding Emiratisation costs. The reform was discussed publicly before the formal announcement, potentially allowing gradual price adjustment. The practical barriers to worker mobility—recruitment debt, information frictions, contractual inertia—may persist despite legal changes. And the concurrent intensification of Emiratisation quotas imposed costs on the low-exposure group that may have masked differential effects on

high-exposure firms.

The paper's contribution is to bound, rather than to estimate, the employer surplus from kafala—while being transparent that the bound applies to the *net* reform package effect and requires assumptions that cannot be fully verified. The bound rules out the large monopsony rents that the system's extreme restrictions might suggest. Whether this means the rents were small all along, were dissipated through other channels, or were capitalized before the event window—these are questions for future research with different data.

Two directions seem particularly promising. First, micro-level evidence on actual wage changes and job-to-job transitions following the reform would test the reform's labor market impact directly, without relying on stock market efficiency in a thin market. Second, studying the reform's effect on worker outcomes—wages, working conditions, recruitment debt—would complete the welfare analysis that this paper begins from the employer side. Legal mobility is worth little to a worker who owes $5,000 to a recruiter. The law changed, but the debt stayed. Quantifying the gap between de jure reform and de facto working conditions would be the most valuable next step.

References

Azar, José, Ioana Marinescu, and Marshall Steinbaum, "Labor Market Concentration," *Journal of Human Resources*, 2022, *57* (S), S167–S199.

Bassier, Ihsaan, Arindrajit Dube, and Suresh Naidu, "Monopsony in Movers: The Elasticity of Labor Supply to Firm Wage Policies," *Journal of Human Resources*, 2022, *57* (S), S50–S86.

Benmelech, Efraim, Nittai Bergman, and Hyunseob Kim, "Strong Employers and Weak Employees: How Does Employer Concentration Affect Wages?," *Journal of Human Resources*, 2022, *57* (S), S200–S250.

Callaway, Brantly and Pedro H.C. Sant'Anna, "Difference-in-Differences with Multiple Time Periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller, "Bootstrap-Based Improvements for Inference with Clustered Errors," *The Review of Economics and Statistics*, 2008, *90* (3), 414–427.

Card, David, Ana Rute Cardoso, Jörg Heining, and Patrick Kline, "Firms and Labor Market Inequality: Evidence and Some Theory," *Journal of Labor Economics*, 2018, *36* (S1), S13–S70.

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer, "The Effect of Minimum Wages on Low-Wage Jobs," *The Quarterly Journal of Economics*, 2019, *134* (3), 1405–1454.

Fama, Eugene F., Lawrence Fisher, Michael C. Jensen, and Richard Roll, "The Adjustment of Stock Prices to New Information," *International Economic Review*, 1969, *10* (1), 1–21.

Fisher, Ronald A., "The Design of Experiments," *Oliver and Boyd*, 1935.

Forstenlechner, Ingo, Fiona Lettice, and Mustafa F. Özbilgin, "Questioning Quotas: Applying a Relational Framework for Diversity Management Practices in the United Arab Emirates," *Human Resource Management Journal*, 2012, *22* (3), 299–315.

International Labour Organization, "Assessment of the Wage Protection System in Qatar," Technical Report, ILO, Geneva 2020.

_ , "UAE Country Profile: Labour Statistics," Technical Report, ILOSTAT 2024.

Kothari, S.P. and Jerold B. Warner, "Econometrics of Event Studies," *Handbook of Empirical Corporate Finance*, 2007, pp. 3–36.

MacKinlay, A. Craig, "Event Studies in Economics and Finance," *Journal of Economic Literature*, 1997, *35* (1), 13–39.

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

Ministry of Human Resources and Emiratisation, "Statistical Report on the Labour Market," Technical Report, MOHRE, Abu Dhabi 2023.

Prager, Elena and Matt Schmitt, "Employer Consolidation and Wages: Evidence from Hospitals," *American Economic Review*, 2021, *111* (2), 397–427.

Robinson, Joan, "The Economics of Imperfect Competition," *Macmillan*, 1933.

Sokolova, Anna and Todd Sorensen, "Monopsony in Labor Markets: A Meta-Analysis," *ILR Review*, 2021, *74* (1), 27–55.

Suresh, Naidu, Nyarko Yaw, and Wang Shing-Yi, "Monopsony Power in Migrant Labor Markets: Evidence from the United Arab Emirates," *Journal of Political Economy*, 2016, *124* (6), 1735–1792.

Toledo, Hugo, "The Political Economy of Emiratization in the UAE," *Journal of Economic Studies*, 2013, *40* (1), 39–53.

UAE Government, "Federal Decree-Law No. 33 of 2021 Concerning the Regulation of Employment Relationship," Official Gazette 2021. Effective February 2, 2022.

UAE Ministry of Human Resources and Emiratisation, "Ministerial Resolution No. 279 of 2022 on Emiratisation," 2022.

A Data Appendix

A.1 Data Sources and Construction

Daily stock price data for all DFM-listed firms were obtained from the Yahoo Finance API using the `yfinance` Python library. The download covered the period January 1, 2019 through December 31, 2024. I retained all firms with at least 500 trading days of data over this period, yielding a sample of 45 firms.

Daily returns were computed as log price changes: $r_{it} = \ln(P_{it}) - \ln(P_{it-1})$, where P_{it} is the adjusted closing price of firm i on day t. Days with zero volume were retained in the sample (they represent trading halts or very illiquid stocks) but were flagged for robustness. The DFM General Index return was computed as the value-weighted average of all sample firm returns, using the previous day's market capitalization as weights.

A.2 Exposure Classification Methodology

The sector classification follows the DFM's official listing categories, which I verified against each firm's annual report and business description. The migrant labor shares used for exposure assignment are sector-level estimates derived from MOHRE aggregate statistics. Specifically, MOHRE reports the number of work permits issued by sector and nationality group, from which I compute the non-national share. These figures are approximate because (a) they cover the entire UAE private sector, not just DFM-listed firms, and (b) some large DFM firms operate across multiple sectors.

To validate the classification, I examined the annual reports of the five largest firms in each exposure group. All high-exposure firms reported workforces consistent with the sector-level migrant shares (e.g., Emaar Properties reports over 90 percent non-national employees; Arabtec—prior to delisting—reported similar figures). All low-exposure firms reported higher national shares, though still majority non-national (e.g., Emirates NBD

reports approximately 40 percent UAE nationals among senior staff).

The continuous migrant share variable is assigned at the sector level as follows: Real Estate, 0.95; Services, 0.90; Industrial, 0.85; Insurance, 0.70; Utilities, 0.70; Banking, 0.65; Financial, 0.65; Telecom, 0.60; Investment, 0.60.

A.3 Event Date Verification

The three event dates were verified against multiple sources: the UAE Government's official portal, MOHRE press releases, legal analyses from international law firms (e.g., Al Tamimi & Company, Baker McKenzie), and contemporaneous media coverage (Gulf News, The National, Khaleej Times). All sources agree on the dates reported in the paper.

I also verified that no major confounding announcements occurred during the $[-1, +3]$ event windows. For Event 1 (September 20, 2021), the main concurrent news was the continuation of post-COVID recovery in UAE real estate—but this would bias *against* a null result by boosting high-exposure firms, not mask a negative effect. For Event 2 (November 15, 2021), no significant macroeconomic announcements were identified. For Event 3 (February 2, 2022), the main concurrent development was the beginning of the Expo 2020 Dubai's final months, which was broadly positive for the UAE economy but not differentially so for high-exposure firms in the expected direction.

B Identification Appendix

B.1 Pre-Trends

The pre-trends test examines whether high-exposure and low-exposure firms followed parallel return paths before the first reform event. In the dynamic DiD specification (Figure 3), I estimate coefficients for each relative trading day in the $[-10, +10]$ window, with $t = -1$ as the reference period. The pre-event coefficients ($t = -10$ through $t = -2$) are jointly

tested against zero. The F-statistic is small and the p-value exceeds 0.50, confirming that the parallel trends assumption is satisfied.

The long-run cumulative returns plot (Figure 7) provides additional visual evidence. Over the two-year pre-reform period (January 2019 through August 2021), the two groups' cumulative returns track each other closely, with temporary divergences that reverse within weeks.

B.2 SUTVA and Spillovers

The SUTVA assumption requires that one firm's reform exposure does not affect another firm's returns. Potential violations include: (a) labor market general equilibrium effects, where workers moving from high-exposure to low-exposure sectors raise wages for the latter; and (b) product market effects, where cost changes in construction affect real estate prices and thus bank balance sheets.

Over the short event windows used here (5 to 11 trading days), these general equilibrium channels are unlikely to operate. Labor reallocation takes months, and product market effects are second-order. SUTVA is more likely violated in the long-run analysis, which I present only as descriptive context and do not interpret causally.

B.3 Randomization Inference Technical Details

The RI procedure permutes the binary exposure indicator across the 45 firms, maintaining the group sizes (18 high, 27 low) in each permutation. For each of the 1,000 permutations, I re-estimate equation (3) and store the coefficient. The two-sided RI p-value is computed as:

$$p_{\text{RI}} = \frac{1}{1000} \sum_{k=1}^{1000} \mathbf{1} \left\{ |\hat{\beta}^{(k)}| \geq |\hat{\beta}^{\text{obs}}| \right\} \tag{7}$$

where $\hat{\beta}^{(k)}$ is the coefficient from permutation k and $\hat{\beta}^{\text{obs}} = 0.0359$ is the observed coefficient. The resulting p-value of 0.354 indicates that 35.4 percent of permuted coefficients exceeded

the observed coefficient in absolute value—a result entirely consistent with the null.

C Robustness Appendix

C.1 Leave-One-Out Analysis

To verify that no single firm drives the main result, I re-estimate equation (3) 45 times, each time dropping one firm from the sample. The coefficient ranges from approximately +2.0 to +5.5 percentage points across the 45 leave-one-out samples, with all p-values exceeding 0.20. No single firm, when removed, changes the qualitative conclusion.

I also conduct a leave-one-sector-out analysis, dropping each of the nine sectors in turn. The coefficient is most sensitive to the exclusion of the real estate sector (the largest high-exposure sector with 7 firms), but remains statistically insignificant in all cases. This confirms that the null result is a feature of the data, not an artifact of one influential observation.

C.2 Alternative Clustering

The main specification clusters standard errors at the firm level. As a robustness check, I also estimate with (a) heteroskedasticity-robust (HC1) standard errors, (b) sector-level clustering, and (c) no clustering (OLS standard errors). The standard errors range from 3.5 percentage points (OLS) to 5.2 percentage points (sector clustered). None of the alternatives produces a significant coefficient.

C.3 Winsorized Returns

To guard against the influence of extreme returns in a thin market, I winsorize daily returns at the 1st and 99th percentiles and re-estimate the main specification. The coefficient changes by less than 0.5 percentage points, confirming that the result is not driven by outlier returns.

C.4 Excluding the Effective Date Event

The effective date (February 2, 2022) is the event with the largest raw CARs for both groups, reflecting a broader market rally. As a robustness check, I re-estimate the main specification using only the first two events (law signing and implementing regulations). The coefficient falls to approximately +1.5 percentage points with a p-value exceeding 0.50. The null result holds whether or not the effective date event is included.

C.5 Continuous Treatment Heterogeneity

The continuous specification (Column 2 of Table 4) uses sector-level migrant share as the treatment intensity. As a further test, I interact the continuous measure with event indicators to test whether the reform's impact varied across the three events. The interaction terms are jointly insignificant, with an F-test p-value exceeding 0.60. This confirms that the null result is not driven by offsetting effects across events.

Acknowledgements

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

Contributors: @ai1scl

First Contributor: `https://github.com/ai1scl`

Project Repository: `https://github.com/SocialCatalystLab/ape-papers`