Private Equity Add-ons, Low Rates, and Local Market Concentration

Dian Xu*

July 25, 2026

Preliminary and Incomplete — Please Do Not Circulate

Abstract

Private equity (PE) add-on acquisitions have become central to U.S. buyout activity, but their competitive effects are hard to measure because targets are often private, small, and local. I build a comprehensive U.S. buy-and-build panel linking 35,118 add-ons, 8,256 platforms, and 2,292 sponsors over 1990–2024, and study their effect on local market outcomes. I find that, first, compared to traditional M&A, PE add-ons are more serial, horizontal, and local. Second, while the effect on national level concentration is muted, PE buy-and-build contribute significantly to higher local market concentration, especially in non-tradable markets and under favorable financial conditions. Third, a counterfactual analysis attributes 1.37 percentage points (34.7 percent) of the 1990–2024 local CR4 rise to PE buy-and-build; the non-tradable contribution is 2.84 percentage points (30.1 percent). Finally, local outcomes show more layoffs, but not broad wage declines or firm and establishment losses.

Keywords: private equity, add-on acquisitions, roll-ups, stealth consolidation, low rates, market concentration

JEL Classification: E52, G24, G34, K21, L11, L41

*London Business School. Email: dxu@london.edu. Dian Xu was supported by the Institute of Entrepreneurship and Private Capital at London Business School during the study.

1 Introduction

> *"In recent years, private equity firms have made serial acquisitions across markets — from nursing homes and apartment buildings to emergency medicine clinics and opioid treatment centres."*
>
> — Lina Khan, Financial Times, *September 21, 2023*

In 2023, an antitrust allegation about private equity (PE) roll-ups in the Texas anesthesia market brought attention to a broader concern about PE serial acquisitions across many local markets (Khan, 2023). The policy concern is that a sequence of individually modest acquisitions can build local market power before any single transaction becomes visible to merger enforcers. This raises the natural question: do PE-backed serial acquisitions increase local market concentration, and do they cause anticompetitive harm in local markets on a large scale?

Private equity add-on acquisitions (or "add-ons" and "bolt-ons") are the subsequent acquisitions made by a platform company after an initial PE buyout. The strategy in which a PE sponsor acquires a platform company and then uses that platform to acquire additional targets is known as a buy-and-build strategy; this is also what "PE serial acquisitions" or "PE roll-ups" typically refer to in the news articles and in policy discussions. I use these terms interchangeably throughout the paper. Add-ons are now central to PE buyout strategies: in 2024, PitchBook records 4,628 completed U.S. add-on Buyout/leveraged buyout (LBO) transactions out of 6,296 completed U.S. Buyout/LBO transactions, or 73.5 percent. In the broader corporate-control benchmark used in this paper, the refined production sample records 4,628 completed U.S. PE add-ons and 11,262 completed non-PE majority-stake control acquisitions in the same year; PE add-ons are therefore 29.1 percent of that broad control-deal count.

Even these counts likely understate the scale of add-on acquisitions because of the under-reporting problem emphasized in the stealth-consolidation literature: after the Hart–Scott–Rodino Antitrust Improvements Act (HSR) reporting threshold rose, newly exempt horizontal mergers increased relative to non-horizontal and above-threshold transactions (Wollmann, 2019); transaction values are often unpublicized (Wollmann, 2023); and disclosure thresholds shape which deals become visible (Barrios and Wollmann, 2024). At the same time, evidence from U.S. administrative microdata shows that mergers can contribute materially to concentration and that national industry screens can miss the relevant margin: in non-tradable local markets, concentration may already breach merger-guideline screens even when national concentration remains below them

(Barkai, Karger, and Loring, 2026).

Low interest rates may also be raising the stakes. A growing literature shows how low interest rates and easy credit can create incentives and opportunities for market power. Liu, Mian, and Sufi (2022) show how low interest rates can raise market power when leaders' investment responds more strongly than followers' investment; Chikis, Goldberg, and López-Salido (2023) qualify that dynamic; and Kroen, Liu, Mian, and Sufi (2024) use high-frequency monetary-policy shocks to show that rate cuts can lower leaders' borrowing costs and increase leader investment and acquisition. PE buy-and-build is a transaction-level setting where this mechanism could operate. A platform acquires a sequence of targets, the sponsor raises capital from limited partners, and add-ons are financed with a mix of fund equity and platform-level debt. Long rates, risky credit spreads, and limited partners' search for yield can therefore affect both deal valuations and the pace of add-on activity.

These considerations raise further questions: do PE buy-and-build strategies combine serial acquisition, product proximity, local geography, and favorable financial conditions in a way that contributes to higher local market concentration?

I answer this by constructing a comprehensive U.S. PE add-on panel from 1990–2024 using PitchBook data and linking it to a U.S. DUNS Marketing Identifier (DMI) county-Standard Industrial Classification (SIC)-year panel of firm-level top-four sales concentration (CR4). A buy-and-build chain begins with a PE sponsor acquiring a platform company—typically through a leveraged buyout—and then directing subsequent add-on acquisitions through that platform. PitchBook records the sponsor's entry and exit dates, the platform's identity, and the identities and dates of each add-on, so I can define each PE-backed chain from the first buyout through exit or, for chains without a recorded exit, a firm-status-corrected close date. The final analysis sample contains 35,118 PE add-ons linked to 8,256 platforms and 2,292 PE sponsors, plus 454,557 non-PE corporate acquisitions for comparison. I match each PE platform and add-on to a county headquarters (HQ) address and map that address into county-SIC markets in the DMI panel. The local DMI panel covers 28.7 million county-SIC-year observations across 1.39 million markets from 1990 to 2024. The treatment surface is deliberately sparse: any PE buy-and-build HQ exposure appears in only 0.123 percent of county-SIC-year markets, and non-PE acquirer-HQ exposure appears in 0.169 percent. The claim is therefore not that PE is everywhere. It is that, in the local markets where PE buy-and-build does appear, concentration changes in a way that is not explained by market fixed effects, state-year shocks, or broad industry-year trends. While the PE literature has documented that add-ons are central to value creation—PE-backed firms acquire faster, add-ons account for a large share of post-

buyout growth, and deal activity responds to favorable financing conditions (Hammer, Knauer, Pflücke, and Schwetzler, 2017; Hammer, Hinrichs, and Schweizer, 2022; Bansraj, Smit, and Volosovych, 2023, 2025)—it has not linked the buy-and-build strategy to local product-market concentration, especially for the private targets that make up the overwhelming majority of add-ons. This paper provides that link.

The evidence begins with three descriptive facts about how PE add-ons differ from non-PE corporate acquisitions. First, PE platforms are more serial acquirers: they are about six times as likely as non-PE corporate acquirers to complete at least three acquisitions and more than twelve times as likely to complete at least ten. This places PE platforms firmly in the serial-acquirer tail of the deal-count distribution and extends the acquisition frequency "boost" documented by Hammer, Knauer, Pflücke, and Schwetzler (2017) and Hammer, Hinrichs, and Schweizer (2022). Second, PE add-ons are more product-close to their acquirers. PE add-ons share the acquirer's PitchBook industry code in 43.4 percent of deals, compared with 18.5 percent for non-PE corporate acquisitions. The embedding cosine similarity between PE add-on targets and their acquirers is 0.596 on average, versus 0.382 for non-PE deals, and the term-frequency–inverse-document-frequency (TF–IDF) cosine gap is equally large. Third, PE add-on targets are closer to their acquirers geographically. PE targets are more likely to share the acquirer's state and five-digit ZIP code (ZIP-5) even though they are less likely to share a county, and after controlling for observable deal characteristics, PE targets are 128 km closer and 1.9 percentage points more likely to share the same ZIP-5. The same pattern holds in the local/non-tradable sub-sample, where PE targets are 249 km closer. Together, these facts imply that PE-backed acquirers buy more targets and that those acquisitions are more horizontal, measured by product similarity and physical distance. Not every add-on is a local roll-up, but PE concentrates more of its deals among product-similar and geographically close acquisitions, especially in markets where local competition is economically meaningful.

Going beyond these descriptive patterns, I ask whether the product and geography proximity translates into horizontal competitive overlap. Antitrust theory identifies product proximity as the key object: horizontal mergers remove close substitutes, and concentration screens are meaningful only inside the relevant product market (Stigler, 1950, 1964; Nocke and Whinston, 2022). Empirical work has made major progress measuring product overlap for public firms—Hoberg and Phillips (2016) show text-based similarity captures competitive relationships missed by SIC and the North American Industry Classification System (NAICS); Fathollahi, Harford, and Klasa (2022) link high-similarity mergers to market-power effects in concentrated industries; and Bonaimé and Wang (2024) provide

price evidence in pharmaceuticals—but PE add-ons overwhelmingly involve private companies. I therefore extend this measurement logic to the private-company universe using PitchBook business descriptions, sentence-transformer embeddings, industry taxonomy, and ZIP-5 geography to build a four-level horizontal classification. H1 is the strictest: a direct competitor in the same local market. H2 keeps direct competition but allows adjacent geography. H3 captures a potential competitor in the same local market. H4 is the broadest: potential competition in an adjacent local market. PE add-ons are classified as H1 in 13 percent of deals and H4 in 44.9 percent, compared with 3.5 percent and 6.7 percent for non-PE acquisitions—a ratio of roughly seven to one at H1 and six to one at H4. Crucially, this gap survives within narrow target-industry by period by state strata in coarsened exact matching (CEM) regressions: PE adds 9.7 percentage points to the H1 probability and 29.2 percentage points to the H4 probability, with the full fixed-effects (FE) and CEM estimates very close to each other. The PE premium in horizontal acquisition is therefore not an industry-composition artifact; within narrow observable strata, PE add-ons are substantially more likely to target direct or potential competitors in the same or adjacent local market.

I next test whether PE add-on activity moves with the financing conditions that determine the PE asset class's cost of capital and deal-execution capacity. In acquirer-year regressions with acquirer and year fixed effects, PE deal activity rises significantly when 10-year Treasury yields are low, when long-rate Federal Open Market Committee (FOMC) surprises ease, and when the Gilchrist and Zakrajšek (2012) excess bond premium (EBP) falls. These are the conditions that reduce the discount rate applied to long-horizon PE fund returns and lower the cost of platform-level leveraged debt, so the direction of the effect is consistent with a financing-cost channel. The evidence on credit spreads is more nuanced. Higher Moody's BAA corporate-bond-yield minus 10-year Treasury spreads (BAA–10Y) or high-yield (HY) spreads are associated with more add-on activity, which seems counterintuitive until one recognizes that high spreads also mean more financially constrained potential targets and therefore more distressed-seller opportunities for PE platforms. The spread level therefore likely captures both the cost of PE financing and the availability of cheap acquisition targets, so its sign in the deal-flow regression does not cleanly identify a financing-cost channel. I do not summarize the evidence as a single low-rate coefficient. The clearest first stage for long-rate and risky-credit financing conditions is from 10-year yields, long-rate monetary-policy surprises, and the Gilchrist–Zakrajsek (GZ) credit-supply measure. Short-rate monetary-policy surprises—which embed policy-news and business-cycle information—are weaker and sometimes negative, consistent with the idea that they move expectations rather than directly reducing long-horizon fi-

nancing costs.

The core concentration result is that PE buy-and-build activity does not robustly increase *national* six-digit-industry concentration, but it does predict higher *local* concentration. The national-level null is unsurprising given that PE targets rarely among the very top firms at the national level. The aggregate-industry concentration literature documents a broad secular rise in national concentration since the 1990s (Grullon, Larkin, and Michaely, 2019; De Loecker, Eeckhout, and Unger, 2020; Kwon, Ma, and Zimmermann, 2024), but that trend is driven predominantly by large publicly traded firms with national market presence. Most PE add-on targets are middle-market private companies—they are not among the national top four in their six-digit industry, so their absorption into a PE platform can leave the national CR4 nearly unchanged. Rossi-Hansberg, Sarte, and Trachter (2021) provide a complementary mechanism: when large multimarket firms simultaneously enter more local geographic markets, their national share rises without any single local market necessarily becoming more concentrated. The PE add-on case is the mirror image—a chain of acquisitions in distinct local geographies can build meaningful local dominance in each market while barely moving the national aggregate. National concentration screens therefore systematically miss the relevant test for a buy-and-build strategy that operates through geographic roll-ups. The local margin is also precisely where stealth consolidation operates: Wollmann (2019) and Wollmann (2023) show that acquisitions below the HSR reporting threshold—exactly the middle-market deals that dominate PE add-ons—are the most likely to escape antitrust scrutiny, and Barkai, Karger, and Loring (2026) document that M&A contributes materially to the observed rise in industry concentration, with local markets capable of breaching merger-guideline screens even when national concentration appears benign.

In the DMI county-SIC-year panel, the baseline PE buy-and-build exposure specification shows that any PE buy-and-build HQ exposure predicts a 0.688 percentage-point higher annual change in local sales-share CR4; platform-HQ exposure predicts 0.900 percentage points; and add-on-HQ exposure predicts 1.174 percentage points. The gradient across concentration measures carries substantial information about the mechanism. The any-PE effect falls from CR4 to CR8 (0.400 percentage points), CR20 (0.168 percentage points), and is near zero at CR50 (0.030). The platform and add-on effects follow the same pattern: platform CR8 is 0.512 and add-on CR8 is 0.793, while both decay to near zero at CR50. This sharp attenuation is precisely what a local roll-up mechanism predicts. Buy-and-build does not compress or restructure the competitive fringe; it concentrates shares at the very top of the local distribution. The platform coefficient exceeding the any-PE coefficient at CR4 (0.900 vs. 0.688) is consistent with the platform company being the or-

ganizational hub of the roll-up: as add-ons are absorbed, the platform accumulates local market share in the markets where it already has a presence. The add-on coefficient being the largest (1.174) reflects a distinct moment in the same process—add-on-HQ exposure marks a target that has just been absorbed, so its local sales share is transitioning toward the acquiring platform's books. Together, the platform and add-on effects identify two nodes in the same consolidation chain: the absorbing entity and the absorbed target, both of which register as local CR4 increases. By contrast, the non-PE corporate acquirer coefficient is -0.058 percentage points and near zero. This asymmetry rules out a generic acquisition story; the concentration effect is specific to the buy-and-build structure, not to acquisitions in general.

The effects are substantially larger in local/non-tradable SIC industries—markets like medical services, food retail, personal services, and local construction, where geographic substitution is limited and a roll-up is most likely to affect the relevant competitive set. Non-tradable industries account for 50.6 percent of DMI county-SIC-year observations and are precisely the markets where stealth consolidation and PE sector-specific concerns are most acute (Gupta, Howell, Yannelis, and Gupta, 2024; Gao, Gupta, and Liu, 2025; Eaton, Howell, and Yannelis, 2020). In those markets, the any-PE effect on the annual CR4 change is 0.754 percentage points, and the platform and add-on effects are 0.985 and 1.860 percentage points, respectively—all substantially larger than the all-industry estimates. The add-on coefficient of 1.860 percentage points in non-tradable markets implies that, per year of PE buy-and-build exposure after acquisition, the relevant local top-four market share rises by nearly two percentage points in the service industries where buy-and-build operates most directly. In a geographically bounded market, a PE platform absorbing a close local rival has no substitute entrant that can offset the concentration effect within the relevant product and geographic scope—the roll-up reshapes the competitive structure of a market that customers cannot easily leave.

The local concentration effect is also sensitive to long-rate and credit conditions, in a way that mirrors the acquisition frequency results. This sensitivity is consistent with how PE financing is structured. Axelson, Jenkinson, Strömberg, and Weisbach (2013) show that PE buyout leverage and pricing are systematically linked to credit market conditions: deals are more levered and priced higher when credit is cheap, suggesting that credit-market booms do not merely shift returns to lenders but also affect which deals happen and at what pace. Add-ons compound this dynamic because each acquisition within an active PE hold period is an additional levered transaction at platform-level debt costs. Long rates, risky credit spreads, and limited partners' search for yield therefore affect not only whether a given buy-and-build program gets started but also how many

add-ons a platform completes before exit. In the concentration regressions, low 10-year yields raise the PE exposure effect on local CR4 by 0.628 percentage points; low BAA–10Y spreads raise it by 0.355 percentage points per year; and low GZ credit-supply spreads raise it by 0.410 percentage points per year. The EBP shock interaction is also positive and significant. In contrast, the monetary-policy surprise interaction is negative and significant. This sign pattern is consistent with the distinction Gertler and Karadi (2015) and Bauer and Swanson (2023) draw between information effects and financial-conditions effects in monetary-policy shocks: high-frequency monetary-policy-surprise (MPS) shocks load on the policy-news component and can move expected future rates in ways that are negatively correlated with current PE activity, while the long-rate and credit-spread measures more directly capture the financing conditions that govern PE leveraged debt and fund valuations. The interest-rate amplification is stronger in local/non-tradable markets, where the low-10Y interaction is 0.649 percentage points per PE buy-and-build exposure year.

To translate the concentration estimates into economic magnitudes, I conduct a counterfactual analysis: for each county-SIC market, I remove PE buy-and-build exposure and recalculate the implied local CR4 using the baseline regression coefficient. It is an accounting exercise that asks: how much of the observed 1990–2024 local CR4 increase can be attributed to the PE channel at current market sizes? The exercise decomposes the PE contribution into two parts—the direct concentration effect of PE buy-and-build exposure (the main coefficient) and the rate-amplification component from the credit-cycle interaction term—and weights each county-SIC-year observation by current sales so that larger markets receive proportionally more weight. The decomposition is informative because the two channels have different policy implications: the main effect is a structural feature of how buy-and-build programs are organized, while the amplification effect is contingent on aggregate credit conditions and therefore moves with the business and credit cycle. For all industries, the sales-weighted local CR4 rose from 82.060 percent at the start of the sample to 86.006 percent by the end—a rise of roughly four percentage points over 1990–2024. Under the counterfactual, any-PE buy-and-build exposure accounts for 1.37 of those percentage points, or 34.7 percent of the total endpoint increase. In the annual breakdown, the direct PE buy-and-build exposure effect contributes 0.0058 percentage points per year; the rate-amplification channel adds a further 0.0333 percentage points per year, implying that roughly 34.7 percent of the all-industry PE contribution operates through the credit-cycle amplification pathway rather than through the direct structural effect. In local/non-tradable markets the magnitudes are substantially larger. The non-tradable local CR4 rose from 67.607 percent to 77.036 percent—a gain of nearly ten percentage

points. PE buy-and-build exposure accounts for 2.84 percentage points of that increase, or 30.1 percent of the endpoint rise—about 2.1 times larger than the all-industry share. In the annual breakdown, the direct PE buy-and-build exposure effect contributes 0.0446 percentage points per year in non-tradable markets; the rate-amplification channel adds a further 0.0365 percentage points per year—roughly equal in magnitude to the direct effect. This near-parity implies that credit-cycle conditions are not merely shifting the timing of PE consolidation but materially affecting its aggregate structural footprint. When long rates are low and credit supply is easy, the same PE buy-and-build footprint produces more concentration, which means more of the observed local CR4 trend is attributable to the PE channel. The PE buy-and-build mechanism therefore connects monetary policy and credit cycles to product-market structure through a channel that operates below the radar of national concentration screens and HSR reporting thresholds.

Finally, I ask whether the same local exposure predicts changes in real economic outcomes, which speaks to whether higher concentration translates into competitive harm. The broader PE buyout literature warns against one-size-fits-all welfare conclusions: classic work emphasizes agency conflicts, leverage, and active monitoring (Jensen, 1989; Kaplan, 1989; Smith, 1990; Lichtenberg and Siegel, 1990); later work decomposes buyout returns into operating gains, multiple expansion, and financing relationships (Guo, Hotchkiss, and Song, 2011; Acharya, Gottschalg, Hahn, and Kehoe, 2013; Ivashina and Kovner, 2011); and establishment-level studies document heterogeneous effects on productivity, employment composition, and operating practices (Boucly, Sraer, and Thesmar, 2011; Davis, Haltiwanger, Handley, Jarmin, Lerner, and Miranda, 2014; Bernstein and Sheen, 2016; Cohn, Nestoriak, and Wardlaw, 2021). The results here are consistent with that heterogeneity. On business dynamics, PE buy-and-build exposure predicts slightly higher establishment counts: the Quarterly Census of Employment and Wages (QCEW) establishment effects are positive in both the full sample and the post-2001 sensitivity, Statistics of U.S. Businesses (SUSB) firm growth is positive, and Business Dynamics Statistics (BDS) establishment entry is higher while exit is lower, though BDS entry and exit estimates are imprecise. Firm deaths in the SUSB panel do not significantly decline. On labor markets, the pattern looks more like restructuring and churn than broad contraction: Quarterly Workforce Indicators (QWI) separations and firm job losses are significantly higher in PE-exposed markets, but firm job gains partly offset the layoff margin; stock employment is close to zero or slightly negative and not statistically significant across QCEW, SUSB, and BDS panels; and wage or stable-earnings effects are inconclusive. The evidence therefore supports a local concentration channel through which rate-sensitive PE buy-and-build activity can reshape market structure, while it

does not support a blanket conclusion that all roll-ups cause broad harm to employment or wages. Whether a particular roll-up is anticompetitive remains a case-by-case question.

The paper makes four contributions. First, it gives the stealth-consolidation literature a scalable local exposure measure: instead of relying on disclosed transaction values or merger filings, it maps PE platform and add-on headquarters into granular county-SIC markets where below-threshold roll-ups are most likely to matter. Second, it contributes to the literature on low interest rates and market power by identifying a specific financial intermediary channel: the financial conditions that support PE deal flow—lower long rates, tighter credit spreads, and credit-supply easing—are the same conditions that amplify the local concentration effect, so the PE buy-and-build channel can transmit monetary-policy and credit-cycle shocks into product-market structure. Third, it disciplines the antitrust interpretation. Sector-specific PE papers show that PE can improve operations in some markets and worsen outcomes in others (Bernstein and Sheen, 2016; Fracassi, Previtero, and Sheen, 2022; Gupta, Howell, Yannelis, and Gupta, 2024; Eaton, Howell, and Yannelis, 2020; Asil, Ramos, Starc, and Wollmann, 2024); this paper estimates the market-structure channel without making a welfare claim, and the outcome evidence on establishments and wages supports caution about blanket conclusions. Fourth, it extends product-proximity measurement to private-company acquisitions by combining PitchBook business descriptions, sentence-transformer embeddings, a four-level horizontal taxonomy, and ZIP-5 geography at the platform-target level, providing a methodology applicable to other contexts where most transactions involve private firms.

2 Data and Sample

2.1 PE Buy-and-Build Sample Construction

The PE sample comes from PitchBook's North America Private Equity universe. I start with completed U.S. add-on acquisitions and retain chains in which the acquiring platform can be linked to a completed Buyout/LBO entry by an identified PE or buyout sponsor. A platform is a portfolio company controlled by one or more PE sponsors; an add-on is a subsequent acquisition by that platform during the observed PE holding window. The construction excludes chains with erratic entry-exit timing and requires at least one add-on on or after platform entry.

PitchBook has a few advantages over other transaction datasets for this question. First, it has broad coverage of U.S. private equity transactions in deal-count terms. Second,

it provides analyst-assigned industry codes for add-on transactions and detailed headquarter location information. Third, it records the sponsor–platform–add-on relation directly, while most merger datasets record the platform buyout and subsequent add-ons as standalone deals with no direct relation. That relation is the object of interest. A buy-and-build strategy is not simply "many acquisitions" by the same legal entity. It is a sponsor-backed organizational design in which a platform is acquired, kept private, and then used as the vehicle for a sequence of follow-on acquisitions. The empirical unit is therefore a chain: sponsor, platform, add-on targets, entry date, and observed exit or end of holding period. The paper uses this chain structure to compare PE-backed acquisition behavior with ordinary corporate mergers and acquisitions (M&A) and then to map the platform/add-on footprint into local markets.

The sample is intentionally conservative. The raw PitchBook add-on flag is broad: it includes transactions linked to platform companies, but not every flagged transaction can be reliably verified as part of a buy-and-build. I therefore require evidence that the platform is PE-backed, that the platform has a PE or buyout sponsor, and that the platform can be placed inside a plausible PE holding window. When multiple PE entries are observed, I use the first entry date for the chain. When the PE exit from the platform is observed in the deal-seller relation, I use the exit deal date to close the holding window. When no clean seller exit is observed but PitchBook's firm-level financing status says the platform is formerly PE-backed, I close the buy-and-build at the last observed add-on date, using the entry date as a floor in the small number of cases where all observed add-ons predate the matched entry. Remaining no-exit chains are treated as ongoing through the fixed sample end. Chains with exits before entries or other inconsistent timing are excluded. This removes some genuine but messy cases, but it keeps the panel auditable and avoids carrying stale PE ownership forward mechanically.

Table 1 reports the filtering cascade. The analysis-window sample contains 35,118 add-ons, 8,256 platforms, and 2,292 sponsors over 1990–2024. The full final platform sample contains 35,623 add-ons and 8,325 platforms. The average disclosed add-on size is about $210.5 million and the median is $32.0 million, but only 8.2 percent of PE add-ons disclose deal value. For this reason, the main empirical analysis uses deal counts, platform/add-on presence, and local-market exposure rather than deal value.

Missing deal values are not surprising, given the private nature of many add-on targets. The stealth-consolidation concern arises precisely because many add-on targets are small, private, and not required to produce the public transaction-value information that would make merger screening straightforward. A value-weighted add-on panel would select toward the most visible and largest transactions. The paper instead uses transaction

counts and market presence as the core exposure measures, then checks whether those exposures matter in local county-SIC markets. This choice sacrifices a clean dollar scale but avoids treating unreported deal values as zero or dropping the transactions most likely to be below ordinary disclosure thresholds.

[Table 1 about here.]

Table 2 then describes the retained PE buy-and-build universe at the deal, sponsor, and platform levels. It shows the difference between the large number of observed add-ons and the much smaller set of disclosed deal values, which is why the paper does not rely on value-weighted acquisition exposure. The table also reports the holding-window audit used for platform presence in the local-market panels.

[Table 2 about here.]

The time-series pattern in the sample also matters for interpretation. The PE add-on strategy became prominent as private-capital supply, regulatory thresholds, and middle-market financing changed. Industry reports describe add-ons as a large share of modern PE buyout count (Bain & Company, 2018; PitchBook Data, Inc., 2025, 2024). The stealth consolidation literature emphasizes the 2001 HSR threshold change (Wollmann, 2019). The post-crisis financing literature emphasizes the shift in middle-market credit intermediation and the expansion of private credit (Davydiuk, Marchuk, and Rosen, 2024; Chernenko, Ialenti, and Scharfstein, 2025). These facts do not identify the paper's regressions by themselves, but they explain why a 1990–2024 PE add-on sample is economically different from a generic public-firm M&A sample. The typical PE platform is a private middle-market company, and the marginal add-on is often a smaller private target rather than a public company large enough to move national concentration statistics.

Figure 1 plots this scale change directly. Panel A compares annual deal counts for PE add-ons and non-PE majority-stake corporate control acquisitions. Panel B shows average disclosed deal size for the same series, making clear both that disclosed add-ons are smaller on average and that deal-value disclosure is sparse for the transactions most relevant to stealth-consolidation concerns.

[Figure 1 about here.]

Figure 2 isolates the PE add-on series itself. The time path shows that the strategy is not simply a static feature of private equity ownership. Add-on volume was small in the early part of the sample, rose after the 2001 Hart–Scott–Rodino notification threshold

revision, and accelerated most sharply after the global financial crisis as private-capital supply and middle-market financing expanded.

[Figure 2 about here.]

The non-PE comparison sample comes from PitchBook's Other North America universe. I keep completed U.S. change-of-control corporate acquisitions, excluding reverse mergers, mergers of equals, and PE-backed transactions. The comparison sample contains 454,557 transactions in the summary-statistics panel. It is not restricted to serial acquirers because the paper's first descriptive question is whether PE ownership itself produces a distinct serial-acquisition tail.

The non-PE benchmark is deliberately broad. Restricting it to serial acquirers would mechanically remove part of the outcome the paper wants to compare. At the same time, the broad benchmark is not a causal control group. Public corporations, private strategic buyers, and family-owned acquirers differ from PE-backed platforms in governance, capital structure, exit horizon, and disclosure. I therefore use the non-PE sample in two roles. In the deal-level sections, it is a descriptive benchmark for serial acquisition, product proximity, and geographic proximity. In the DMI regressions, non-PE M&A exposure is a diagnostic exposure and placebo-style comparison, not a matched counterfactual for PE treatment.

Table 3 reports the corresponding summary statistics for this broad non-PE corporate M&A benchmark. The table is meant to show the scale and deal-value disclosure pattern of the comparison sample, not to define a matched control group.

[Table 3 about here.]

2.2 Product, Geography, and Industry Mapping

PitchBook provides proprietary industry groups, sectors, and industry codes. For product-market similarity, I use sentence-transformer embeddings of PitchBook business descriptions. The embedding cosine is validated against shared PitchBook industry codes and achieves a receiver operating characteristic area under the curve (ROC AUC) of 0.783, outperforming term-frequency–inverse-document-frequency (TF–IDF) cosine and keyword-overlap benchmarks. Appendix A.3 gives the full validation.

For national concentration tables, I map platforms to six-digit NAICS using S&P Capital IQ name, alias, former-name, and website-domain matches, then SIC-to-NAICS conversion and an empirical PitchBook-industry to NAICS crosswalk. The platform-level match rate is 99.7 percent. Appendix A.4 reports the pass-by-pass procedure.

For DMI local concentration, the market is county by SIC-4 by year. I map PitchBook platform and add-on headquarters ZIP codes to counties and map PitchBook companies to SIC-4 using PitchBook SIC relation files and crosswalked industry-code fallbacks. Appendix A.5 documents this PitchBook-to-SIC layer. This assignment is intentionally HQ based: it is the consistently observed geography for both PE and non-PE transactions. It is not a plant-level ownership transfer.

2.3 DMI County-SIC-Year Local Concentration Panel

The local concentration backbone is the Dun & Bradstreet (D&B) U.S. DUNS Marketing Identifier (DMI) establishment panel. I form county-SIC-year markets, group establishments into firms using headquarters Data Universal Numbering System identifiers (HQ-DUNS) when available and DUNS-level singletons otherwise, and compute firm-level sales-share CR4 for each county-SIC-year. The dependent variable in the main local regressions is the annual change in CR4 in percentage points. The regression panel keeps consecutive county-SIC years with positive lagged employment and sales. It covers 28,723,791 county-SIC-year rows and 1,386,329 markets over 1990–2024.

Government establishments are excluded from the regression sample: county-SIC markets with SIC major group 91–97 (public administration) are excluded from the regression panels. The tradable/non-tradable split follows the DMI SIC coverage audit: construction (SIC 15–17), retail (52–59), real estate (65), and services (70–89) are local/non-tradable, subject to the paper's regression exclusions of SIC 82, SIC 86, public administration 91–97, and commercial banking SIC 6011. The retained non-tradable sample contains 14.55 million CR4 regression rows, or 50.6 percent of the saved CR4 panel.

Table 4 summarizes the DMI county-SIC-year regression surface used in the local concentration analysis. It lists the market definition, sample size, concentration outcome, weighting convention, and the rarity of PE and non-PE M&A exposure in the regression panel.

[Table 4 about here.]

The DMI establishment data are tilted toward local services, which is also the main area of interest for PE roll-ups and stealth consolidation. In the raw DMI establishment-year data, local/non-tradable industries account for 65.8 percent of entries, while tradable production industries account for 16.0 percent of rows. The remainder is made up of SIC observations classified as ambiguous finance/holding or unknown/nonclassifiable. Sales weights paint a different picture because tradable markets are sales-heavy: local/non-

tradable rows account for 36.5 percent of sales, while tradable production rows account for 47.3 percent. This is why the paper separates unweighted regression estimates from sales-weighted descriptive aggregation and counterfactual accounting.

3 PE Add-On Characteristics

This section studies the characteristics of PE add-ons, and how they differ from traditional corporate acquisitions. Section 3.1 documents the descriptive facts around the scale, serial acquisition behavior, and industry and geographic proximity. Section 3.2 then uses text-based product similarity and nested product-geography classifications to test whether PE add-ons are more horizontal in a product-market sense. The final subsection studies whether acquisition activity is sensitive to long-rate and credit conditions.

3.1 Descriptive Facts

The first descriptive fact is scale. I document that PE add-ons represent a substantial current share of U.S. corporate control activity. In 2024, PitchBook records 4,628 completed PE add-ons, compared with 11,262 completed non-PE majority-stake control acquisitions in the broad benchmark used in Figure 1. Thus PE add-ons account for 29.1 percent of current control acquisition deal count. In other words, roughly one in every three corporate control acquisitions is a PE add-on. Thus, add-ons are not just common inside PE buyouts, they have become a large part of the corporate-control market.

Buy-and-builds typically involve a series of add-ons. In the PE sample, the platform add-on distribution has a long right tail. The distinction from ordinary corporate M&A is clearest in the threshold distribution. Table 5 shows that PE platforms are much more likely than non-PE corporate acquirers to cross repeated-acquisition thresholds. At three and ten acquisitions, PE platforms sit far deeper in the serial-acquirer tail.

[Table 5 about here.]

This fact aligns with Hammer, Hinrichs, and Schweizer (2022), who document that PE-backed portfolio companies acquire more after PE ownership. Lifetime acquisition count is informative on a different dimension: it shows that in addition to a simple boost effect, PE ownership also pushes acquirers into the serial-acquirer tail. This is important because repeated horizontal acquisitions, while individually small, can accumulate substantial market power within a segmented market.

Simple industry and geographic proximity. The second descriptive fact is that PE add-ons are closer to their platforms on simple product and geography measures. The distinction matters because antitrust markets, which define where firms compete, have both a product dimension and a geographic dimension. A same-industry deal between two national manufacturers is not the same competitive event as an add-on between two outpatient clinics in the same neighborhood. The raw same-code and distance patterns therefore motivate, but do not by themselves constitute, the text-based product-market result below.

Table 6 reports the raw proximity differences. PE add-ons are more likely than non-PE corporate acquisitions to share the same PitchBook industry code (43.4 versus 18.5 percent) and the same broader industry group (63.3 versus 31.2 percent). The business description text measures point in the same direction. Term-frequency–inverse-document-frequency (TF–IDF) cosine and sentence-transformer embedding cosine are both substantially higher for PE add-ons. Geography is more nuanced: PE add-ons are slightly less likely to be in the same state or same county, but they are more likely to share the same five-digit ZIP code (ZIP-5) and are 162 km closer on average. This combination is important. PE roll-ups do not simply buy the nearest target in the same state; they often appear to trade off geographic expansion and tight local-market consolidation within a narrow product space. The county boundary gives useful discipline. Mapping target and acquirer ZIP-5 codes to counties in the 1990–2024 analysis-window sample shows that 10.5 percent of PE add-on targets are in the same county as the platform, compared with 12.1 percent of non-PE acquisition targets, or 0.87 times the non-PE rate. Thus the PE geographic premium is not a broad same-county premium; it is concentrated at the exact-ZIP tail and in the shorter continuous-distance distribution. Panel D repeats the same raw comparison inside target industries classified as local/non-tradable, where geography is most directly connected to competitive overlap.

[Table 6 about here.]

3.2 Text-Based Product-Market Proximity

Going beyond the descriptive proximity measures, I construct a private-company taxonomy of horizontal acquisitions. The goal is to move beyond the observation that PE deals often share broad industry codes and ask whether they are more horizontal in a product-market sense. Methodologically, I encode platform and target business descriptions into sentence-transformer vectors, compute pairwise product similarity, and combine that continuous similarity measure with industry-code consistency and ZIP/county-based geog-

raphy rules to form the nested H1–H4 product-market classes used below. Appendix A.3 and Appendix A.8 provide the full implementation, validation, and threshold details.

The classification has four nested stringency layers. H1 is the strictest layer: a direct competitor in the same local market. H2 keeps the direct-competitor requirement but allows adjacent-market geography. H3 relaxes product overlap to potential competitors in the same local market. H4 is the broadest layer, allowing potential competitors in adjacent markets. The construction is intentionally conservative in non-tradable industries because geography is part of the market definition there; for tradable industries, a product-market competitor can matter even without close physical distance.

Table 7 shows that the PE horizontal premium appears at every layer. At H1, 13 percent of PE add-ons qualify as direct competitors in the same market, compared with 3.5 percent of non-PE acquisitions. At H4, 44.9 percent of PE add-ons qualify, compared with 6.7 percent of non-PE acquisitions. The PE–non-PE gap therefore grows as the definition expands from the narrowest direct-local competitor to the broader potential-adjacent competitor. This pattern is hard to explain as a pure artifact of industry mix, because the largest gap occurs when the definition explicitly combines product and market proximity.

[Table 7 about here.]

Selection remains a concern: PE sponsors may concentrate in industries where all acquisitions are more horizontal. I therefore estimate controlled linear probability models with increasingly saturated fixed effects. The richest column uses target-industry group by five-year period by target-state strata, equivalent to a coarsened exact matching (CEM) comparison in which the PE effect is identified only from strata containing both PE and non-PE deals. In these controlled models, the PE coefficient is 9.7 percentage points at H1 and 29.2 percentage points at H4. The coefficients are smaller than the raw differences, which means composition matters, but they remain large and highly significant. The residual PE premium is the deal-side evidence that PE add-ons are systematically closer in product space after controlling for observable industry, time, and geography.

[Table 8 about here.]

As a companion descriptive diagnostic, geography shows the same pattern after controls. Conditional on year, industry group, state, and the matched industry-period-state strata, PE targets are about $|-128|$ km closer to their acquirers than non-PE targets and 1.9 percentage points more likely to share a ZIP-5. The same-county and same-state coefficients are also positive in the most saturated specifications, but both are smaller and statistically imprecise; the core proximity evidence is the shorter-distance and same-ZIP-5

pattern. The interpretation is not that all PE add-ons are local. It is that PE add-ons have a fatter local tail than non-PE acquisitions, exactly the tail that matters for county-level service-market concentration.

[Table 9 about here.]

The tradability split sharpens the interpretation. Non-tradable goods and services must be consumed locally, so geographic proximity is a precondition for competitive overlap. In the non-tradable subsample, PE targets are $|-249|$ km closer to their acquirers than non-PE targets; in tradable industries, the distance gap is $|-96|$ km. The same-ZIP-5 coefficient is positive in both groups, while the same-county coefficients are small and imprecise. The broader same-state row is 6.3 percentage points for non-tradables but only 0.6 percentage points for tradables, reinforcing that the economically relevant proximity premium is concentrated where local competition is meaningful.

[Table 10 about here.]

Figure 3 visualizes the same tradability split by plotting the full target–acquirer distance distribution. The figure shows that the non-tradable PE distance distribution has more mass close to the acquirer, which is the part of the distribution most relevant for local service-market competition.

[Figure 3 about here.]

The final proximity fact is that product and geographic proximity are not the same object. Table 11 shows that the correlation between embedding cosine and distance is small for PE add-ons. High-product and close-geography deals account for 25.8 percent of PE add-ons, while high-product and far-geography deals account for 24.2 percent. This near-balance suggests two PE rollout modes. One mode consolidates within a local market; the other replicates a standardized platform model across geographies while staying close in product space. Non-PE acquisitions exhibit a stronger correlation between product and geographic proximity, consistent with a more integrated adjacency strategy.

[Table 11 about here.]

Together, the descriptive facts and text-based classifications establish the central measurement premise. PE add-ons are not merely numerous; they are systematically closer to the acquirer in product space and, in the markets where geography matters, in physical space. The evidence also explains why national industry measures are too coarse for the

paper's main question. A platform that buys a clinic, dental practice, HVAC contractor, veterinary clinic, or insurance broker may not change national CR4, but it can remove a close rival from a county-level service market. That is why the outcome analysis uses DMI county-SIC markets rather than only national NAICS aggregates.

3.3 Sensitivity to Long Rates and Credit Conditions

The rate channel follows from the financing structure of buy-and-build. PE funds raise committed capital from limited partners, acquire a platform, and finance add-ons with a mixture of fund equity and platform-level debt. Lower long-term rates reduce discount rates and fund hurdle rates; compressed credit spreads lower the cost of deal debt; and search-for-yield can increase the capital allocated to PE. Short-rate and monetary policy shocks are more subtle because they also arrive with business-cycle information.

The acquirer-year evidence asks whether PE-backed acquirers' annual deal counts are more sensitive to aggregate financial conditions than non-PE corporate acquirers. The baseline estimating equation is

$$Y_{k,t} = \beta \left(PE_{k,t} \times Z_t \right) + \mathrm{FE}_k + \mathrm{FE}_t + \varepsilon_{k,t}, \tag{1}$$

where $Y_{k,t}$ is the number of acquisitions completed by acquirer k in year t, $PE_{k,t}$ indicates that the acquirer is PE-backed after PE entry, Z_t is the standardized annual rate, spread, or monetary-policy-shock measure, FE_k and FE_t are acquirer and year fixed effects, and $\varepsilon_{k,t}$ is the regression error. The coefficient β is therefore the incremental PE sensitivity to that aggregate condition; the PE main effect is absorbed by acquirer fixed effects.

Table 12 reports the full baseline rate-channel regression table. It is the paper's main acquirer-year regression table for long rates, credit spreads, and the joint rate-spread specification, with industry-fixed-effect columns shown beside the preferred acquirer-fixed-effect columns.

[Table 12 about here.]

Table 13 then gives a compact summary for the broader set of financial-condition measures used in the local-concentration analysis. The "Add-on frequency response" column reports $\hat{\beta}$ from equation (1) for a one-standard-deviation increase in the row variable. Negative yield, spread, EBP, or term-spread coefficients therefore mean PE deal frequency rises when the listed condition is lower. The MPS rows are different. They report raw positive short-rate surprises, and the orthogonalized MPS response is larger

than the raw MPS response (0.0989 versus 0.0677). Because purging Greenbook and pre-FOMC macro news does not mute the coefficient, I treat these rows as cautionary shock diagnostics rather than as financing-cost evidence or as the low-rate sign convention used in the local-concentration interactions below.

[Table 13 about here.]

The credit-spread rows should therefore be read as deal-frequency evidence on a different margin from the local-concentration table below. In the acquirer-year count panel, spread-level and term-spread responses are negative, consistent with PE deal frequency being higher in lower-spread and lower-term-spread years. Table 18, by contrast, signs the local CR4 interactions as lower spreads, a lower term spread, or easier credit and asks whether PE-exposed local markets experience larger CR4 increases when financing conditions ease.

4 PE Buy-and-Build and Market Concentration

The paper's concentration claim is local, not national. The national Economic Census tables are useful context because they show what the paper should not claim: PE add-ons do not robustly increase national six-digit-industry concentration. The DMI county-SIC-year panel is the main outcome surface because it can observe narrow product-geography markets where roll-ups are economically plausible. The empirical logic therefore runs in two steps. First, the paper shows that a PE add-on target is usually too small to move national concentration ratios in a six-digit industry. Second, it asks whether the same activity predicts higher concentration growth inside county-SIC markets after absorbing market fixed effects, state-year shocks, and broad industry-year shocks.

4.1 National Concentration Effects

This subsection tests whether PE buy-and-build activity predicts changes in national industry concentration using a pure first-difference design at the six-digit NAICS industry level. The test is designed before the local DMI analysis, and it establishes a benchmark: if PE roll-ups primarily consolidate narrow product-geography niches, they should leave national sales-share concentration largely unchanged.

Panel construction. The estimation sample is built from U.S. Economic Census industry concentration tables for the 2012, 2017, and 2022 waves, yielding two five-year inter-

vals (2012–2017, 2017–2022) and 1,322 regression observations for the concentration-ratio columns. Industries are harmonized to strict one-to-one 2017 NAICS codes to avoid spurious concentration changes driven by reclassification. Treatment counts are built by matching PitchBook add-on deals and platform entries to the NAICS-6 industry of the acquirer, then aggregating the cumulative count over each five-year interval. The resulting panel contains 35,441 matched add-on deals and 8,240 matched platform entries.

The outcome variables are five-year changes in the national sales-share concentration ratios $\Delta C_{i,t}^{(n)}$ for $n \in \{4, 8, 20, 50\}$ and the Herfindahl–Hirschman Index (HHI). Because the 2012 Census API did not publish HHI, the HHI column uses only the 2017–2022 interval (814 observations). Level counts are used for the treatment rather than log-one-plus because many NAICS-6 industries have zero PE activity in a given interval, and the log transform would drop those zero-treatment industries from the identifying variation.

Estimating equation. The regression is

$$\Delta C_{i,t}^{(n)} = \beta \, T_{i,t-5\to t} + \gamma \, C_{i,t-5}^{(n)} + \mathrm{FE}_t + \varepsilon_{i,t}, \tag{2}$$

where i is a six-digit NAICS industry, $t \in \{2017, 2022\}$, $\Delta C_{i,t}^{(n)}$ is the five-year change in the national sales share of the top n firms, $T_{i,t-5\to t}$ is the cumulative level count of PE add-ons or platform entries in the interval, and $C_{i,t-5}^{(n)}$ is lagged concentration. The coefficient β is the treatment-count slope, γ is the lagged-concentration coefficient, FE_t denotes period fixed effects, and $\varepsilon_{i,t}$ is the regression error. The lagged concentration term controls for mean reversion, while period fixed effects absorb any common trend across census waves. Regressions are weighted by lagged industry sales to reflect economic scale, and standard errors are clustered at the NAICS-3 level.

Results. Table 14 reports the estimates for both treatment panels.

For Panel A (add-on counts), the CR4 coefficient is -0.002 ($t = -1.36$), the CR8 coefficient is -0.001 ($t = -0.69$), and the CR20 and CR50 coefficients are similarly small and insignificant. For Panel B (platform entries), the CR4 coefficient is -0.027 ($t = -1.91$), with the remaining ratios also negative but insignificant. Across both panels and all four concentration ratios, there is no evidence that PE buy-and-build activity increases national industry concentration.

The HHI results look different: both the add-on count (-0.12) and the platform-entry count (-0.98) are negative and highly significant. This pattern is not contradictory. HHI equals the sum of squared national market shares, so it is sensitive to the arrival of a new

middle-market entrant in an otherwise thinly-populated tail of the size distribution. A PE-backed platform that acquires a series of small niche targets does not move the top-four share, but it does add a new firm group of moderate size, mechanically reducing HHI. The negative HHI coefficient is therefore consistent with PE activity adding new mid-tier competitors at the national industry level rather than consolidating the existing large players.

The lagged concentration coefficients are negative and significant for CR4 in both panels (-0.024 for add-on counts and -0.027 for platform entries), which reflects standard mean reversion in industry concentration: industries with higher concentration at the prior census wave tend to see smaller (or negative) changes over the next five years, independent of PE activity.

Collectively, these results say that PE buy-and-build strategies do not manifest as increases in national six-digit-industry concentration ratios. This is theoretically expected. PE roll-up targets are typically private middle-market firms with no measurable share of a broad national industry. The same roll-up that consolidates an orthodontic practice cluster in one metro area or a regional HVAC service network leaves the national CR4 virtually unchanged. The absence of a national effect is therefore diagnostic rather than disappointing: it confirms that PE roll-ups operate at a finer product-geography grain than the Economic Census can measure, which directly motivates the DMI county-SIC-year panel used in the next subsection.

[Table 14 about here.]

4.2 Local Concentration Effects

The main DMI pure-exposure specification is

$$\begin{aligned} \Delta C_{m,t}^{(k),sales} = \beta_e PE_{m,t}^{e} + X'_{m,t-1}\Gamma \\ + \mathrm{FE}_m + \mathrm{FE}_{s(m),t} + \mathrm{FE}_{sic2(m),t} + \varepsilon_{m,t}, \qquad k \in \{4, 8, 20, 50\}, \end{aligned} \tag{3}$$

where m is a county-SIC-4 market, t is year, e indexes the PE exposure definition, and $k \in \{4, 8, 20, 50\}$ indexes the concentration ratio. $C_{m,t}^{(k),sales}$ is the firm-level sales-share concentration ratio for the top k local firm groups, $PE_{m,t}^{e}$ is the indicated PE buy-and-build exposure, $X_{m,t-1}$ is the vector of lagged controls, Γ is the corresponding control-coefficient vector, β_e is the exposure coefficient, and $\varepsilon_{m,t}$ is the regression error. The terms FE_m, $\mathrm{FE}_{s(m),t}$, and $\mathrm{FE}_{sic2(m),t}$ are market, state-by-year, and SIC-2-by-year fixed effects, respectively. The controls include the outcome-specific lagged concentration ratio, lagged

establishments, employment, sales, lagged PE platform counts, PE add-on counts, non-PE M&A deal counts, and non-PE M&A exposure. Regressions are unweighted county-SIC-year market regressions; sales weights are used only for descriptive aggregation and counterfactual accounting. Financial-condition interactions are reported separately in Section 4.3.

The exposure variable is deliberately based on PE buy-and-build HQ exposure rather than only disclosed add-on counts. This choice addresses two measurement problems. First, below-threshold or undisclosed targets can make pure add-on counts an understatement of the local buy-and-build footprint, the same concern emphasized by the stealth-consolidation literature. Second, the PE proximity evidence shows that add-ons are often local to the platform, so platform-HQ and add-on-HQ exposure are useful reduced-form exposure surfaces even when exact plant ownership links are incomplete. In this design, the exposure is not a literal plant-level treatment. It is a granular marker for where the PE buy-and-build footprint is plausibly active.

The rarity of treated markets is also informative. In the CR4 regression sample, any PE buy-and-build HQ exposure appears in only 0.123 percent of county-SIC-year markets, platform-HQ exposure in 0.095 percent, and add-on-HQ exposure in 0.034 percent. A stricter serial-acquirer audit on all non-empty DMI county-SIC-year markets gives a similar message: the local roll-up footprint is sparse by market count, and it overlaps little with the broad non-PE serial-acquirer footprint. Among 27.0 million non-empty DMI county-SIC-year markets from 1996–2025, 0.0747 percent have PE serial-platform exposure, 0.1097 percent have non-PE serial-acquirer exposure, 0.0068 percent have both, and 99.8224 percent have neither.

[Table 15 about here.]

Figure 4 shows the geography behind the same exposure surface. The map is a descriptive diagnostic rather than an identification assumption: it shows that platform-HQ and add-on-HQ exposure is sparse in the full county-SIC panel but expands across major metropolitan corridors between 2000 and 2024.

[Figure 4 about here.]

These facts help with interpretation, but they do not make the design experimental. The coefficients are identified from within-market changes after market, state-year, and SIC-2-year absorption; they are not randomized treatment effects. The identification argument is more modest: the exposure surface is extremely granular, broad state-industry-year shocks are absorbed, local market size and lagged concentration are controlled, and

the main effects are stronger for PE than for non-PE acquisition exposure and stronger in non-tradable markets. That combination is what supports the paper's local concentration channel.

Table 16 reports the unconditional local concentration regressions. The columns mirror the national concentration table by showing CR4, CR8, CR20, and CR50 outcomes side by side, while the panels rotate through the three HQ-based PE exposure definitions.

[Table 16 about here.]

The pure-exposure results support the local concentration channel. Any PE buy-and-build HQ exposure predicts a 0.688 percentage-point larger annual CR4 change, platform-HQ exposure predicts 0.900 percentage points, and add-on-HQ exposure predicts 1.174 percentage points. The same pattern remains positive and statistically significant through CR20. The CR50 column is smaller, which is consistent with PE exposure moving the leading local ownership groups more than the entire top-fifty fringe. The estimates are economically meaningful despite the low exposed-market share because exposed markets are not randomly chosen tiny markets. They are markets where a PE platform or add-on target is headquartered, and those markets are disproportionately sales-relevant in the local DMI surface. This is why the paper keeps the regressions unweighted, to avoid letting already large markets dominate identification, but uses sales weights for descriptive aggregation and counterfactual accounting.

Why using CR4 instead of HHI. The baseline outcome is firm-level sales-share CR4 rather than HHI. This is a measurement choice. DMI HQ-DUNS links identify many firm groups, but some subsidiary establishments have missing HQ identifiers. Those missing links can make establishment-level HHI too noisy because the denominator contains many singleton establishments that may in fact be commonly owned. CR4 is less sensitive to that long tail and focuses on the leading local firm groups. The ownership diagnostics support this choice. In the latest DMI audit, 75.2 percent of rows marked as subsidiaries lack an HQ DUNS, so the missing-HQ problem is real. But those no-HQ subsidiary rows account for only 2.5 percent of full local-market sales, and no-HQ code-0 groups account for only 6.7 percent of top-four sales. HQ-DUNS groups account for 80.5 percent of full local-market sales and 80.2 percent of top-50 sales. Thus, the problematic missing-subsidiary mass is too small in sales terms to overturn the top-four concentration measure, while dropping all missing-HQ rows would incorrectly delete part of the competitive fringe. HHI remains useful as a robustness surface, but it is not yet a validated owner-group HHI.

The weighting choice also matters. Rossi-Hansberg, Sarte, and Trachter (2021) show that local and national concentration can move differently. In this paper, descriptive aggregation weights local markets by sales rather than employment because the concentration ratio is a sales-share object. A DMI weighting audit illustrates why this matters: sales-weighted local concentration is 84.7 percent for CR4 and 93.3 percent for CR50, while employment-weighted CR4 is much lower. The gap is mechanical and economically relevant. Sales-heavy local markets account for 72.6 percent of DMI sales but only 5.6 percent of employment. Employment weighting would therefore mute precisely the markets in which a sales-share concentration outcome is largest. Sales weighting also avoids muting concentration gains in markets where rising market power is associated with lower labor share, as emphasized by Barkai (2020).

4.3 Rate and Credit Condition Effects

The local CR4 rate evidence is strongest for long-rate and credit-spread proxies. Any-PE interactions are positive for low 10-year yields, low 5-year yields, low 10Y–2Y term spreads, low Moody's BAA corporate-bond-yield minus 10-year Treasury spreads (BAA–10Y), low Gilchrist–Zakrajsek (GZ) spreads, low excess bond premium (EBP), and low high-yield option-adjusted spreads (HY OAS). The 2-year yield row is weaker, and the displayed orthogonalized monetary-policy-surprise (MPS) row is negative. This pattern is consistent with the deal-side rate evidence in Section 3.3: long risky financing and credit-spread conditions are the relevant PE buy-and-build states, while high-frequency short-rate surprises are a different policy-shock object. Table 13 reinforces that caution because the orthogonalized MPS coefficient is larger than the raw MPS coefficient in the acquirer-year deal-frequency design. The MPS family therefore should not be rationalized simply as raw macro-news contamination or used as financing-cost evidence. This is why the paper uses "rate-sensitive" rather than "low-rate causal effect" language. The same PE exposure can load positively on slow-moving long-rate and credit-spread states and negatively on short-rate or high-frequency policy-shock states without a coding contradiction.

The financial-condition variables are continuous, standardized indexes, not binary low-rate indicators. I use standardized versions rather than raw levels because the rows combine yields, credit spreads, EBP, and FOMC-window surprises with different units and sample variances. Let $z_r(X_t)$ denote the standardization of annual raw variable X_t

over the nonmissing years in row r:

$$z_r(X_t) = \frac{X_t - \mu_r(X)}{\sigma_r(X)}, \qquad Ease_t^X = -z_r(X_t). \tag{4}$$

Thus, the baseline low-10Y variable is $Low10Y_t = -z(\overline{DGS10}_t)$, where $\overline{DGS10}_t$ is the annual mean daily 10-year Treasury yield. The two MPS variables are first aggregated within year t and then standardized using the same row-specific transformation. The coefficient on $PE_{m,t} \times Ease_t^X$ is therefore the additional percentage-point change in local CR4 associated with a one-standard-deviation easing in proxy X, conditional on PE buy-and-build exposure. Table 19 gives the full construction, source, and economic interpretation for each proxy.

The credit-spread sign differs from the acquirer-year frequency table because the two designs ask different questions. Table 13 reports PE add-on frequency in the high-spread direction for the spread rows; a positive response can reflect valuation or distress opportunities when financially constrained targets become easier to buy. Table 18 signs the same family of conditions as low spreads/easier credit and asks whether PE-exposed local markets see larger CR4 increases. Distressed target availability can raise deal frequency in high-spread years without implying that local concentration effects are largest then, because tight credit also constrains financing, execution, and post-acquisition roll-up scale.

For auditability, Table 17 reports the full low-10Y CR4 interaction regressions, including the control stack, t-statistics, fixed-effect indicators, and market cluster counts.

[Table 17 about here.]

Table 18 collects the any-PE local CR4 interactions across the main financial-condition proxies. It shows which long-rate, term-spread, credit-spread, and MPS rows support the rate-sensitive concentration channel and which rows discipline the interpretation. The interaction coefficient in each row is in percentage points of annual local CR4 change per one-standard-deviation easing in that row's proxy.

[Table 18 about here.]

The PE main-effect column should be read row by row, not as a single unconditional PE coefficient. Each row estimates a separate regression with a different financial-condition interaction and, in some cases, a different available sample window. For the low-10Y row, the main effect is the PE effect when the standardized low-10Y variable is at

its sample mean; the concentration effect in that specification is carried primarily by the PE-by-low-10Y interaction. The other rows condition on different rate or credit proxies, so their PE main effects need not match the first row.

The HHI robustness surface supports part, but not all, of the CR4 rate story. Establishment-sales-share HHI gives positive any-PE main effects and positive interactions for low credit spreads and low 10Y–2Y term spreads, but it does not support the low-10Y-yield row. This is useful discipline. The preferred low-10Y CR4 table is the right input for the paper's sales-share CR4 counterfactual, but the broader rate evidence should be read as proxy-specific. A stronger causal statement would require a pre-specified long-rate or credit supply shock that moves PE acquisition incentives without directly changing local product-market demand.

4.4 Non-Tradable Markets and Non-PE Serial Acquirer Comparison

The local-market interpretation predicts stronger effects in non-tradable industries, where geographic segmentation makes local roll-ups more likely to affect competition. The retained non-tradable CR4 sample contains 14.55 million rows, or 50.6 percent of the saved CR4 panel, after excluding public administration, commercial banking, schools, membership organizations, and other markets outside the local-service definition used in the paper. Table 20 confirms the predicted pattern in the analogous low-10Y CR4 design. In non-tradable SIC markets, the any-PE main effect is 0.754 percentage points, compared with 0.088 in the full-sample low-10Y CR4 specification. The main table uses the combined any-PE buy-and-build presence indicator, which equals one when either the platform HQ or an add-on target HQ is present in the county-SIC-year. The platform-HQ and add-on-HQ split is reported separately in Appendix A.6.

[Table 19 about here.]

The comparison with non-PE acquisition exposure is informative. In the low-10Y CR4 specification, non-PE acquirer exposure has essentially no main effect (-0.058), while the any-PE main effect is 0.088. Any-PE and non-PE M&A exposure overlap in only 5,895 county-SIC-year markets, which is 16.7 percent of any-PE markets, 12.2 percent of non-PE markets, and 7.6 percent of markets with either exposure. The low-10Y interaction is positive for both PE and non-PE acquisition exposure, which means the rate component should not be sold as exclusively PE-specific. The stronger PE main effect, the larger non-tradable effect, the limited overlap, and the horizontal/geographic targeting evidence together make PE add-ons the more plausible local concentration mechanism. This is

an acquisition and ownership-channel interpretation, not a generic claim that all low-rate business-dynamics effects are PE-driven. Low-rate leader-investment mechanisms may operate in the background, but in this design they are auxiliary: the exposure is tied to PE platform/add-on M&A presence, and the outcome is local firm-group sales-share CR4. The establishment, employment, and business dynamics panels below are therefore descriptive discipline around the concentration evidence, not replacements for the acquisition channel.

4.5 Counterfactual Analysis

The counterfactual is a mechanical accounting exercise, not a structural merger simulation. For each year, I weight county-SIC markets by current DMI market sales and apply the all-industry or non-tradable CR4 coefficients:

$$\Delta C_t^{PE} = \sum_m w_{m,t} \left[\widehat{\beta} PE_{m,t} + \widehat{\theta} (PE_{m,t} \times Low10Y_t) \right]. \tag{5}$$

The no-PE path subtracts the cumulative contribution from the observed sales-weighted CR4 level path. It does not model entry, prices, quality, ownership restructuring, or general-equilibrium responses. In Equation (5), ΔC_t^{PE} is the PE accounting contribution in year t, $w_{m,t}$ is the current-sales weight for county-SIC market m, $\widehat{\beta}$ is the estimated PE main-effect coefficient, $\widehat{\theta}$ is the estimated low-10Y interaction coefficient, $PE_{m,t}$ is the any-PE buy-and-build exposure indicator, and $Low10Y_t$ is the standardized low-10-year-yield condition. Current sales weights are the headline because they match the observed CR4 level path used in the descriptive DMI concentration figures. The weights are normalized over all county-SIC markets in the relevant sample-year, including markets with no PE footprint, so the reported 1.37 and 2.84 percentage-point totals are sample-wide contributions to aggregate local CR4 rather than effects conditional on PE-exposed markets. The underlying coefficients are local-market treatment increments for exposed county-SIC-year markets, conditional on the regression controls and fixed effects; the counterfactual total is the sales-weighted aggregation of those increments after untreated markets receive zero mechanical PE contribution. Lagged sales weights, using one-year-lagged market sales, and equal county-SIC weights are retained as diagnostics rather than as the main economic weighting convention.

Figure 5 shows the time-series version of the counterfactual accounting exercise. The top panels compare the observed sales-weighted CR4 level with the mechanical no-PE path, while the bottom panels decompose the cumulative PE contribution into the main

exposure component and the low-10Y amplification component.

[Figure 5 about here.]

[Figure 6 about here.]

Table 21 summarizes the same accounting exercise as annual average and cumulative percentage-point contributions. It compares all industries with local/non-tradable industries and reports the any-PE, platform-HQ, and add-on-HQ exposure definitions separately.

[Table 20 about here.]

Table 22 answers the corresponding conditional question by re-normalizing current-sales weights only within county-SIC-year markets with any PE buy-and-build exposure. The affected set is tiny by unweighted market counts but much larger in sales-weighted terms: all-industry affected markets account for 0.116% of county-SIC-year observations and 5.8% of current DMI sales weight. The conditional contribution is therefore large, but the sample-wide aggregate is not obtained by dividing by the unweighted market share. Year by year, it is the affected-market contribution multiplied by the affected markets' current-sales weight in the full sample.

[Table 21 about here.]

For Any PE buy-and-build HQ exposure, the sample-wide all-industry contribution is 1.369 percentage points over 1990–2024. Local/non-tradable industries contribute 2.841 sample-wide percentage points, about 2.1 times the all-industry magnitude. Relative to observed endpoint changes, the Any-PE contribution equals 34.7 percent of the all-industry sales-weighted CR4 increase and 30.1 percent of the local/non-tradable increase. The denominators are endpoint changes in the observed sales-weighted CR4 level series: all-industry CR4 rises from 82.060 to 86.006, while local/non-tradable CR4 rises from 67.607 to 77.036. The implied no-PE endpoint changes remain positive, about 2.58 percentage points in the all-industry sample and 6.59 percentage points in the local/non-tradable sample. Thus, PE exposure is not the whole rise in local concentration. It is a large accounting component of the observed increase under the preferred reduced-form CR4 coefficients.

4.6 Local Economic Outcomes

This subsection asks what happens to local real-side outcomes around PE buy-and-build activity. The exercise is deliberately not a mechanism test in the narrow sense, and it is not a concentration regression: the public and administrative panels below observe establishments, firms, employment, wages, worker flows, and business dynamics, but they do not observe local ownership shares, owner-level CR4, prices, or quality. The goal is therefore descriptive discipline around the local concentration evidence: does the same exposure surface coincide with entry, exit, establishment growth, layoffs, employment changes, and wages?

Table 23 reports the source-native local economic outcome surface. The Quarterly Census of Employment and Wages (QCEW) is the closest county-NAICS6 local-service companion to DMI for establishments, employment, and wages. Statistics of U.S. Businesses (SUSB) adds true employer-firm, establishment, and employment stock outcomes at state-NAICS6. Quarterly Workforce Indicators (QWI) adds county-NAICS4 labor flows, and Business Dynamics Statistics (BDS) adds county-sector business dynamics. I do not treat these as four reporting schemes from one administrative microdata file. Their administrative bases overlap but differ: QCEW is a BLS establishment panel built primarily from state unemployment-insurance records; QWI is a Census Longitudinal Employer-Household Dynamics product built from state unemployment-insurance wage records; SUSB is a Census Business Register tabulation; and BDS is compiled from the Census Longitudinal Business Database, which links Business Register establishments and firms over time.[1] Agreement across these panels therefore triangulates across related employer-data systems rather than replicating the same underlying file at different reporting schemes. I checked the public SUSB files used in the production pipeline: the state-NAICS6 and county-NAICS3 tables report stock variables–firms, establishments, employment, payroll, and, in the 2022 county file, receipts– but not flow variables such as firm entry, firm births, firm deaths, establishment entry, or establishment exit. Those entry/death margins are therefore available here only from BDS, whose county-sector panel is much coarser than the local product markets where PE roll-ups operate. For readability, the QCEW rows in Table 23 report only the full 1990Q2–2024Q4 window. Appendix A.11 reports the 2001Q1–2024Q4 support-stable sensitivity, which drops the thinner BLS reconstructed 1990–2000 county-NAICS6 support and yields very similar baseline magnitudes.

[1] Official BLS/Census documentation describes QCEW as administrative data from state unemployment-insurance programs; QWI as workforce composition and turnover measures calculated from unemployment-insurance wage records; SUSB as data extracted from the Census Business Register; and BDS as tabulations compiled from the Longitudinal Business Database, which is constructed by linking annual Business Register snapshots.

[Table 22 about here.]

The firm-count margin points up rather than down. PE buy-and-build exposure is associated with higher SUSB firm growth, and the BDS firm-stock row is also positive. On the establishment margin, QCEW establishment growth is positive in the full sample, with the appendix 2001+ sensitivity showing the same economic pattern. BDS establishment entry is higher and establishment exit is lower, although neither estimate is statistically significant. Because the BDS panel collapses local activity to county-sector units, narrow local effects are likely averaged into broad sector units; the imprecision should therefore be read as a coarse-data limitation, not as evidence that local establishment dynamics are absent. BDS firm deaths also fall, but the estimate is imprecise.

The labor-market reading is churn rather than a clean employment-collapse story. QWI separations and firm job losses are significantly higher in exposed markets. QWI firm job gains are also positive, however, so the layoff margin is partly offset by job gains. Stock employment is not significantly lower in a robust way: the SUSB, BDS, and QCEW employment rows are close to zero or slightly negative, while the QWI employment row is modestly positive, and none of these employment-stock estimates is statistically significant. The wage evidence is also inconclusive, with positive but imprecise QCEW wage-per-worker effects and an imprecise QWI stable-earnings row. Appendix A.11 reports the detailed QCEW local economic outcome magnitudes for establishments, employment, and wages in the full 1990Q2–2024Q4 sample and the 2001Q1–2024Q4 support-stable sensitivity.

The QCEW results are deliberately modest. Establishment growth is positive in both the full sample and the 2001+ sensitivity: the P90 treated effect is 0.81 percent in the full sample and 0.78 percent in the 2001+ sensitivity. Employment effects are close to zero (-0.01 percent and -0.01 percent at the same comparison), and wage-per-employee effects are positive but imprecise (0.49 percent and 0.58 percent). The local outcome evidence therefore does not support a broad claim that PE add-ons destroy firms, establishments, or employment everywhere. It is more consistent with modest firm and establishment growth alongside labor market reallocation. QCEW has no firm boundary, owner identifier, price, or quality measure, so welfare claims still require price, quality, and worker-level evidence at the roll-up and local-market level.

5 Conclusion

PE add-ons have become a central form of U.S. acquisition activity. This paper shows why they matter for market structure. PE-backed acquirers are unusually serial, their targets are more horizontal and geographically proximate than comparable non-PE acquisitions, and their deal flow rises when long rates and risky credit conditions are favorable. These deal-level facts make PE buy-and-build a plausible financial-intermediation channel linking low rates to local consolidation.

The concentration evidence is local and strongest where local market definition is most economically meaningful. National Economic Census regressions show little evidence that PE add-ons move national industry concentration, which is consistent with middle-market private targets being too small to change six-digit national CR4. In DMI county-SIC markets, however, PE exposure predicts larger increases in firm-level sales-share CR4, and the estimates are larger in non-tradable industries where customers, workers, and competitors are more geographically segmented. The rate interactions also require careful interpretation. The amplification evidence is strongest for low long rates and favorable risky-credit conditions, whereas high-frequency MPS shocks are weaker or negative in the local concentration design and do not map cleanly to the long-horizon financing conditions relevant for PE. A mechanical counterfactual implies that PE exposure accounts for a meaningful share of the observed 1990–2024 local CR4 increase, especially when markets are weighted by sales, but this accounting exercise should be read as a market-structure magnitude rather than a structural welfare estimate.

The paper does not claim that every roll-up is anticompetitive. QCEW establishment, employment, and wage results are mixed rather than uniformly negative. The right implication is narrower and more useful: serial PE acquisitions can build local concentration in ways that national statistics and per-deal thresholds miss, and the intensity of that channel changes with long-rate and credit conditions. Antitrust review should therefore evaluate cumulative acquisition patterns in local product-geography markets, while macro-finance research on low rates and market power should treat PE-sponsored buy-and-build as a specific acquisition channel rather than an abstract financing wedge.

References

Acharya, V. V., O. F. Gottschalg, M. Hahn, and C. Kehoe. Corporate Governance and Value Creation: Evidence from Private Equity. *Review of Financial Studies*, 26(2):368–402, 2013. doi: 10.1093/rfs/hhs117.

Asil, A., P. Ramos, A. Starc, and T. G. Wollmann. Painful Bargaining: Evidence from Anesthesia Rollups. NBER Working Paper No. 33217, November, 2024.

Axelson, U., T. Jenkinson, P. Strömberg, and M. S. Weisbach. Borrow Cheap, Buy High? The Determinants of Leverage and Pricing in Buyouts. *Journal of Finance*, 68(6):2223–2267, 2013. doi: 10.1111/jofi.12082.

Bain & Company. Global Private Equity Report 2018. Industry report, Bain & Company, 2018.

Bansraj, D. S., H. T. Smit, and V. Volosovych. Buy-and-Build Strategies. Working paper, 2023.

Bansraj, D. S., H. T. Smit, and V. Volosovych. Buy-and-Build Strategies. Working paper, Erasmus University Rotterdam, 2025.

Barkai, S. Declining Labor and Capital Shares. *Journal of Finance*, 75(5):2421–2463, 2020. doi: 10.1111/jofi.12909.

Barkai, S., E. Karger, and S. Loring. M&A and the Rise in Industry Concentration. Working paper, Boston College / Federal Reserve Bank of Chicago / Counterfactual Consulting; January revision, 2026.

Barrios, J. M., and T. G. Wollmann. A New Era of Midnight Mergers: Antitrust Risk and Investor Disclosures. *American Economic Journal: Microeconomics*, 16(4):77–111, 2024. doi: 10.1257/mic.20220064.

Bauer, M. D., and E. T. Swanson. An Alternative Explanation for the "Fed Information Effect". *American Economic Review*, 113(3):664–700, 2023.

Bernstein, S., and A. Sheen. The Operational Consequences of Private Equity Buyouts: Evidence from the Restaurant Industry. *Review of Financial Studies*, 29(9):2387–2418, 2016. doi: 10.1093/rfs/hhw037.

Bonaimé, A., and Y. Wang. Mergers, Product Prices, and Innovation: Evidence from the Pharmaceutical Industry. *Journal of Finance*, 79(3):2195–2236, 2024. doi: 10.1111/jofi.13321.

Boucly, Q., D. Sraer, and D. Thesmar. Growth LBOs. *Journal of Financial Economics*, 102(2):432–453, 2011. doi: 10.1016/j.jfineco.2011.05.014.

Chernenko, S., R. Ialenti, and D. Scharfstein. Bank Capital and the Growth of Private Credit. Working paper, 2024 draft, 2025.

Chikis, C. A., J. Goldberg, and D. López-Salido. A Comment on "Low Interest Rates, Market Power, and Productivity Growth". *Econometrica*, 91(6):2457–2461, 2023. doi: 10.3982/ECTA20621.

Cohn, J., N. Nestoriak, and M. Wardlaw. Private Equity Buyouts and Workplace Safety. *Review of Financial Studies*, 34(10):4832–4875, 2021. doi: 10.1093/rfs/hhab001.

Davis, S. J., J. Haltiwanger, K. Handley, R. Jarmin, J. Lerner, and J. Miranda. Private Equity, Jobs, and Productivity. *American Economic Review*, 104(12):3956–3990, 2014. doi: 10.1257/aer.104.12.3956.

Davydiuk, T., T. Marchuk, and S. Rosen. Direct Lenders in the U.S. Middle Market. *Journal of Financial Economics*, 162:103946, 2024. doi: 10.1016/j.jfineco.2024.103946.

De Loecker, J., J. Eeckhout, and G. Unger. The Rise of Market Power and the Macroeconomic Implications. *Quarterly Journal of Economics*, 135(2):561–644, 2020. doi: 10.1093/qje/qjz041.

Eaton, C., S. T. Howell, and C. Yannelis. When Investor Incentives and Consumer Interests Diverge: Private Equity in Higher Education. *Review of Financial Studies*, 33(9): 4024–4060, 2020. doi: 10.1093/rfs/hhz129.

Fathollahi, M., J. Harford, and S. Klasa. Anticompetitive Effects of Horizontal Acquisitions: The Impact of Within-Industry Product Similarity. *Journal of Financial Economics*, 144(2):645–669, 2022. doi: 10.1016/j.jfineco.2021.06.017.

Fracassi, C., A. Previtero, and A. Sheen. Barbarians at the Store? Private Equity, Products, and Consumers. *Journal of Finance*, 77(3):1439–1488, 2022. doi: 10.1111/jofi.13134.

Gao, J., A. Gupta, and T. Liu. Private equity and the hospital industry. Working paper, available at SSRN, 2025.

Gertler, M., and P. Karadi. Monetary Policy Surprises, Credit Costs, and Economic Activity. *American Economic Journal: Macroeconomics*, 7(1):44–76, 2015. doi: 10.1257/mac.20130329.

Gilchrist, S., and E. Zakrajšek. Credit Spreads and Business Cycle Fluctuations. *American Economic Review*, 102(4):1692–1720, 2012. doi: 10.1257/aer.102.4.1692.

Grullon, G., Y. Larkin, and R. Michaely. Are US Industries Becoming More Concentrated? *Review of Finance*, 23(4):697–743, 2019. doi: 10.1093/rof/rfz007.

Guo, S., E. S. Hotchkiss, and W. Song. Do Buyouts (Still) Create Value? *Journal of Finance*, 66(2):479–517, 2011. doi: 10.1111/j.1540-6261.2010.01640.x.

Gupta, A., S. T. Howell, C. Yannelis, and A. Gupta. Owner Incentives and Performance in Healthcare: Private Equity Investment in Nursing Homes. *Review of Financial Studies*, 37(4):1029–1077, 2024. doi: 10.1093/rfs/hhad082.

Hammer, B., A. Knauer, M. Pflücke, and B. Schwetzler. Inorganic Growth Strategies and the Evolution of the Private Equity Business Model. *Journal of Corporate Finance*, 45: 31–63, 2017. doi: 10.1016/j.jcorpfin.2017.04.006.

Hammer, B., H. Hinrichs, and D. Schweizer. What Is Different about Private Equity-Backed Acquirers? *Review of Financial Economics*, 40(2):117–149, 2022. doi: 10.1002/rfe.1128.

Hoberg, G., and G. Phillips. Text-Based Network Industries and Endogenous Product Differentiation. *Journal of Political Economy*, 124(5):1423–1465, 2016.

Ivashina, V., and A. Kovner. The private equity advantage: Leveraged buyout firms and relationship banking. *Review of Financial Studies*, 24(7):2462–2498, 2011.

Jensen, M. C. Eclipse of the Public Corporation. *Harvard Business Review*, 67(5):61–74, 1989.

Kaplan, S. N. The Effects of Management Buyouts on Operating Performance and Value. *Journal of Financial Economics*, 24(2):217–254, 1989. doi: 10.1016/0304-405X(89)90047-0.

Khan, L. It's time to halt roll-up schemes that violate antitrust laws. Financial Times, September 2023. URL `https://www.ft.com/content/93103af9-768a-4545-9166-20389c254edc`. Opinion essay by the chair of the Federal Trade Commission.

Kroen, T., E. Liu, A. Mian, and A. Sufi. Falling Rates and Rising Superstars. Working paper, January revision, 2024.

Kwon, S. Y., Y. Ma, and K. Zimmermann. 100 Years of Rising Corporate Concentration. *American Economic Review*, 114(7):2111–2140, 2024. doi: 10.1257/aer.20220621.

Lichtenberg, F. R., and D. Siegel. The Effects of Leveraged Buyouts on Productivity and Related Aspects of Firm Behavior. *Journal of Financial Economics*, 27(1):165–194, 1990.

Liu, E., A. Mian, and A. Sufi. Low Interest Rates, Market Power, and Productivity Growth. *Econometrica*, 90(1):193–221, 2022. doi: 10.3982/ECTA17408.

Nocke, V., and M. D. Whinston. Concentration Thresholds for Horizontal Mergers. *American Economic Review*, 112(6):1915–1948, 2022. doi: 10.1257/aer.20201038.

PitchBook Data, Inc. Q3 2024 US PE Middle Market Report. Industry report, PitchBook Data, Inc., 2024. Published December; sponsored by Morningstar DBRS.

PitchBook Data, Inc. 2024 Annual US PE Breakdown. Industry report, PitchBook Data, Inc., 2025. Published January.

Rossi-Hansberg, E., P.-D. Sarte, and N. Trachter. Diverging Trends in National and Local Concentration. *NBER Macroeconomics Annual*, 35(1):115–150, 2021. doi: 10.1086/712317.

Smith, A. J. Corporate Ownership Structure and Performance: The Case of Management Buyouts. *Journal of Financial Economics*, 27(1):143–164, 1990.

Stigler, G. J. Monopoly and Oligopoly by Merger. *American Economic Review: Papers and Proceedings*, 40(2):23–34, 1950.

Stigler, G. J. A Theory of Oligopoly. *Journal of Political Economy*, 72(1):44–61, 1964.

Wollmann, T. G. Stealth Consolidation: Evidence from an Amendment to the Hart-Scott-Rodino Act. *American Economic Review: Insights*, 1(1):77–94, 2019. doi: 10.1257/aeri.20180137.

Wollmann, T. G. Terms of the Deal Were Not Announced: Accounting for Unpublicized M&A. *AEA Papers and Proceedings*, 113:284–287, 2023. doi: 10.1257/pandp.20231099.

Table 1: PE Add-On Sample: Filtering Cascade

Step	Restriction	Add-ons	% base	Platforms	Sponsors
1	Completed U.S. add-ons	57,967	100.0%	15,762	—
2	PE-backed/corporate platform type	53,827	92.9%	15,383	3,223
3	Identifiable sponsor entry buyout/LBO	37,538	64.8%	9,122	—
4	Plausible holding window	35,623	61.5%	8,325	2,309
5	Analysis window, 1990–2024	35,118	60.6%	8,256	2,292

Note: The retention column uses completed U.S. add-ons as the base. Step 2 excludes missing platform links (3,314 add-ons), unmatched type cases (35 platforms / 36 add-ons), and known non-PE/corporate platform types (344 platforms / 790 add-ons). Step 3 requires a Buyout/LBO entry tied to an identified PE/buyout sponsor (19,016 entry-deal rows) and matches 8,290 seller-exit rows. Step 4 drops 797 chains / 1,915 add-ons with implausible holding-window timing. The component flags are 28 chains / 127 add-ons with exit-before-entry timing and 780 chains / 1,819 add-ons with no observed add-on on or after matched PE entry; categories can overlap. For platform-presence windows, 1,811 formerly PE-backed no-exit platforms close at the last observed add-on date; 3,111 remain ongoing, reducing analysis-window platform-years by 21,558.

Table 2: Summary Statistics: PE Buy-and-Build Universe

Measure	**Value**
Panel A: Sample counts, analysis window	
Add-on deals	35,118
Platforms with add-ons	8,256
PE/buyout sponsors	2,292
Sponsor-restricted entry-deal rows	18,756
Matched seller-exit rows	8,195
Panel B: Deal terms and industry-code overlap	
Add-on deal-size disclosure	8.2%
Mean disclosed add-on size ($M)	210.5
Median disclosed add-on size ($M)	32.0
Median entry stake acquired	100.0%
Median exit stake sold	100.0%
Same PitchBook industry code	43.4%
Panel C: Platform holding-window audit	
Platforms with observed seller exit	3,403
Missing exits closed by former-PE status	1,811
Platforms still treated as ongoing	3,111
Long no-exit chains before correction	1,909
Long no-exit chains after correction	452
Platform-years under old no-exit rule	84,053
Platform-years under corrected rule	62,495
Platform-year reduction	21,558 (25.6%)
Mean holding window (yrs)	7.1
Median holding window (yrs)	5.6
Panel D: Serial-acquisition tail	
Mean add-ons per platform	4.3
Median add-ons per platform	2.0
Maximum add-ons by one platform	455
Platforms with $\geq$10 add-ons	740
Platforms with $\geq$50 add-ons	45
Sponsors with $\geq$10 add-ons	969
Sponsors with $\geq$50 add-ons	269

Note: PitchBook U.S. PE buy-and-build sample, 1990–2024. Add-on deal size is in USD millions and is reported only for disclosed transactions. Seller exits are matched through PitchBook's deal-seller relation. Missing seller exits are closed at the last observed add-on date only when the platform's firm-level financing status is "Formerly PE-Backed"; otherwise the platform is treated as ongoing. Same PitchBook industry code is a descriptive target–platform code match, not the horizontal-classification measure introduced below. Platform-year counts refer to the 1990–2024 analysis window used for local-market exposure.

Table 3: Summary Statistics: Non-PE Corporate M&A Comparison

Measure	**Value**
Panel A: Sample counts, analysis window	
Corporate M&A transactions	454,557
Unique targets	324,305
Unique acquirers	32,565
Acquirers with $\geq$5 deals	5,348
Acquirers with $\geq$10 deals	2,383
Panel B: Deal terms and classification	
Deal-size disclosure	44.5%
Mean disclosed deal size ($M)	267.2
Median disclosed deal size ($M)	12.2
P75 disclosed deal size ($M)	143.0
Horizontal transactions (both-side)	21.5%
Panel C: Serial-acquirer activity	
Mean deals per acquirer	3.6
Median deals per acquirer	1.0
P75 deals per acquirer	3.0
Maximum deals by one acquirer	449

Note: PitchBook Other North America completed U.S. corporate M&A sample, 1990–2024, excluding PE-backed transactions, reverse mergers, and mergers of equals. Deal size is in USD millions and is reported only for disclosed transactions.

Table 4: DMI County-SIC-Year CR4 Regression Sample

Metric	Value
Market	County $\times$ SIC-4 $\times$ year
Regression period	1990–2024
Rows	28,723,791
Markets	1,386,329
Regression weights	None
Sales-weighted mean CR4	84.5
Employment-weighted mean CR4	55.1
Any PE B&B HQ market-year share	0.123%
Platform HQ market-year share	0.095%
Add-on HQ market-year share	0.034%
Non-PE M&A market-year share	0.169%

Note: DMI county-SIC-year CR4 panel. The dependent variable is the annual change in firm-level sales CR4, in percentage points. Firms are HQ-DUNS groups when available and DUNS-level singletons otherwise. Exposure shares are computed in the regression sample.

Table 5: Serial Acquirer Comparison: PE-Backed vs. Non-PE Corporate

$\geq n$ **deals**	**PE Count**	**PE %**	**Corp Count**	**Corp %**
≥ 2	5,050	60.7%	69,306	20.3%
≥ 3	3,446	41.4%	23,971	7.0%
≥ 4	2,494	30.0%	12,774	3.7%
≥ 5	1,909	22.9%	8,271	2.4%
≥ 10	748	9.0%	2,288	0.7%
≥ 20	231	2.8%	458	0.1%
≥ 50	45	0.5%	21	0.0%

Note: PE sample: PitchBook U.S. PE-backed platforms with at least one add-on. Non-PE sample: PitchBook Other North America corporate acquirers, 1990–2024. The table reports the share of acquirers exceeding each acquisition-count threshold.

Table 6: Geographic and Product Proximity: PE Add-On vs. Non-PE Acquisitions

Measure	PE Mean	Non-PE Mean	Diff	t-stat	N_{PE}	$N_{\text{Non-PE}}$
Panel A: Geographic Proximity						
Same State	0.215	0.240	-0.025^{***}	-10.39	34,732	207,761
Same County	0.105	0.121	-0.016^{***}	-8.30	30,511	170,495
Same ZIP-5	0.058	0.047	$+0.011^{***}$	7.86	32,037	186,159
Distance (km)	1,258	1,421	-162^{***}	-22.70	31,959	184,641
Panel B: Product Proximity						
Same Industry Code	0.434	0.185	$+0.249^{***}$	90.46	35,598	220,622
Same Industry Group	0.633	0.312	$+0.321^{***}$	117.40	35,598	220,622
Panel C: Text-Based Proximity						
TF-IDF Cosine	0.220	0.110	$+0.110^{***}$	130.50	35,576	220,552
Embedding Cosine	0.596	0.382	$+0.214^{***}$	240.88	35,576	220,552
Panel D: Non-Tradable Target Industries						
Same State	0.240	0.204	$+0.036^{***}$	6.70	7,545	33,695
Same County	0.087	0.098	-0.012^{***}	-3.04	6,782	28,596
Same ZIP-5	0.046	0.043	$+0.003$	1.02	7,054	31,315
Distance (km)	1,119	1,393	-274^{***}	-18.88	7,018	31,076
Same Industry Code	0.506	0.219	$+0.287^{***}$	46.88	7,656	35,399
Same Industry Group	0.696	0.334	$+0.362^{***}$	62.23	7,656	35,399
TF-IDF Cosine	0.250	0.120	$+0.130^{***}$	66.99	7,653	35,446
Embedding Cosine	0.632	0.398	$+0.233^{***}$	123.34	7,653	35,446

Notes: Full sample: 35,623 PE add-on and 488,336 non-PE corporate acquisitions from PitchBook (1990–2024). Panel D restricts to target PitchBook industry groups classified as local/non-tradable (7,660 PE and 111,032 non-PE deals before row-specific missingness). Each measure is computed on the subsample with valid data for both target and acquirer; N columns report the actual observation count. Stars denote significance of the PE–non-PE difference (Welch's t-test for unequal variances): *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Geographic measures are binary (1 = same location) except Distance (haversine great-circle km between target and acquirer ZIP-5 centroids, via `pgeocode`/ZCTA); Same County uses the project ZIP-county crosswalk. Both-sides ZIP-5 coverage: 89.9% PE, 38.1% non-PE. Product measures use PitchBook analyst-assigned industry codes and groups. Text-based measures require both-sides descriptions (99.9% PE, 45.2% non-PE); TF-IDF uses 10,000 text features and Embedding Cosine uses all-MiniLM-L6-v2 sentence embeddings.

Table 7: Horizontal Acquisition Rates: Main Classification (H1, H2, H3, H4)

Level	Definition	PE	Non-PE	Diff	t	N_{PE}
H1	Direct competitor, same market	14.5%	1.9%	+12.6%***	66.5	34,981
H2	Direct competitor, adjacent market	16.0%	2.0%	+14.0%***	71.1	34,981
H3	Potential competitor, same market	41.1%	6.4%	+34.7%***	130.4	34,981
H4	Potential competitor, adjacent market	44.9%	6.7%	+38.1%***	141.9	34,981

Notes: Sample restricted to deals with valid descriptions for both target and acquirer (99.9% of PE, 45.2% of non-PE). Product proximity = embedding cosine similarity (all-MiniLM-L6-v2): direct tiers use top quintile (embed $\geq$ 0.73) and potential tiers use above median (embed $\geq$ 0.60). Direct tiers also require same business role; potential tiers do not. Roles are parsed from the first word of each PitchBook description and mapped to service (Provider/Operator/Owner), production (Manufacturer/Developer/Producer/Designer/Publisher), or distribution (Distributor/Supplier/Wholesaler/Retailer). "Same market" = target distance $\leq$ 288 km (PE Q1) for non-tradable industries; no distance requirement for tradable. "Adjacent market" = distance $\leq$ 1080 km (PE median). Diff is the raw PE–non-PE rate difference, equivalently the PE coefficient from an unadjusted LPM with an intercept and PE indicator, not the controlled LPM on the next slide. Stars use Welch t: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 8: PE Effect on Horizontal Classification: LPM with Controls

	(1)	(2)	(3)	(4)	(5)
Panel A: H1					
PE-backed	0.1072***	0.1081***	0.0985***	0.0980***	0.0967***
	(0.0160)	(0.0155)	(0.0139)	(0.0137)	(0.0140)
Non-PE mean	0.0352	0.0352	0.0352	0.0352	0.0352
N	245,789	245,774	245,774	245,774	242,756
Panel B: H2					
PE-backed	0.1204***	0.1210***	0.1099***	0.1092***	0.1071***
	(0.0116)	(0.0112)	(0.0112)	(0.0111)	(0.0115)
Non-PE mean	0.0371	0.0371	0.0371	0.0371	0.0371
N	245,789	245,774	245,774	245,774	242,756
Panel C: H3					
PE-backed	0.2844***	0.2911***	0.2731***	0.2715***	0.2678***
	(0.0347)	(0.0339)	(0.0275)	(0.0272)	(0.0273)
Non-PE mean	0.1199	0.1199	0.1199	0.1199	0.1199
N	245,789	245,774	245,774	245,774	242,756
Panel D: H4					
PE-backed	0.3154***	0.3215***	0.2996***	0.2975***	0.2919***
	(0.0183)	(0.0178)	(0.0166)	(0.0166)	(0.0171)
Non-PE mean	0.1264	0.1264	0.1264	0.1264	0.1264
N	245,789	245,774	245,774	245,774	242,756
Year FE	No	Yes	Yes	Yes	Yes
Industry group FE	No	No	Yes	Yes	Yes
State FE	No	No	No	Yes	Yes
CEM (Ind×Year×State)	No	No	No	No	Yes

Notes: Linear probability model estimates. Dependent variable in each panel is a binary indicator for whether the deal qualifies as horizontal at that classification level. Sample restricted to deals with both-sides company descriptions and valid embedding cosine similarity. Column (5) uses three-way interaction fixed effects (target industry group × 5-year period × target state), equivalent to coarsened exact matching: only deals in strata containing both PE and non-PE observations contribute to identification. Indeterminate classifications (missing role or distance data) treated as non-horizontal (conservative). Standard errors clustered at the target industry group level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 9: PE Effect on Geographic Proximity: OLS/LPM with Controls

	(1)	(2)	(3)	(4)	(5)
Panel A: Distance (km)					
PE-backed	-156.8***	-185.8***	-157.1***	-127.9***	-127.8***
	(45.3)	(43.2)	(35.2)	(31.1)	(31.0)
Non-PE mean	1415	1415	1415	1415	1415
N	207,904	207,889	207,889	207,889	204,984
Panel B: Same ZIP-5 code					
PE-backed	0.0108*	0.0142**	0.0185***	0.0191***	0.0192***
	(0.0063)	(0.0061)	(0.0060)	(0.0059)	(0.0058)
Non-PE mean	0.0473	0.0473	0.0473	0.0473	0.0473
N	207,904	207,889	207,889	207,889	204,984
Panel C: Same county					
PE-backed	-0.0155*	-0.0102	0.0004	0.0070	0.0066
	(0.0086)	(0.0084)	(0.0095)	(0.0089)	(0.0090)
Non-PE mean	0.1203	0.1203	0.1203	0.1203	0.1203
N	192,835	192,817	192,817	192,817	190,019
Panel D: Same state					
PE-backed	-0.0183	-0.0111	0.0052	0.0181	0.0197
	(0.0134)	(0.0133)	(0.0134)	(0.0123)	(0.0120)
Non-PE mean	0.2331	0.2331	0.2331	0.2331	0.2331
N	207,904	207,889	207,889	207,889	204,984
Year FE	No	Yes	Yes	Yes	Yes
Industry group FE	No	No	Yes	Yes	Yes
State FE	No	No	No	Yes	Yes
CEM (Ind×Year×State)	No	No	No	No	Yes

Notes: Panel A: OLS with distance in km as dependent variable. Panels B–D: linear probability models with binary same-ZIP-5, same-county, and same-state indicators. Sample restricted to deals with valid target and acquirer ZIP-5 codes (distance computable), valid deal year, and valid target industry group. The same-county panel additionally requires valid target and acquirer county assignments from the project ZIP-county crosswalk. Column (5) uses three-way interaction fixed effects (target industry group × 5-year period × target state), equivalent to coarsened exact matching. Standard errors clustered at the target industry group level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 10: PE Geographic Proximity Effect: Tradability Heterogeneity

	Non-tradable				Tradable			
	Distance	Same ZIP-5	Same county	Same state	Distance	Same ZIP-5	Same county	Same state
PE-backed	-249.0***	0.0105	0.0051	0.0626***	-95.8***	0.0216***	0.0076	0.0064
	(31.1)	(0.0058)	(0.0072)	(0.0117)	(24.5)	(0.0075)	(0.0113)	(0.0112)
Non-PE mean	1391	0.0435	0.0976	0.1974	1420	0.0481	0.1249	0.2403
N	36,637	36,637	34,028	36,637	171,228	171,228	158,769	171,228
Year + Industry + State FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Subsample regressions with year + target industry group + target state fixed effects. Non-tradable = target in PitchBook industry groups classified as local (healthcare services, restaurants, retail, etc.). Same-county columns require valid target and acquirer county assignments from the project ZIP-county crosswalk. Standard errors clustered at the target industry group level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 11: Joint Distribution of Product and Geographic Proximity

	PE Add-Ons	Non-PE
Panel A: Correlation		
Pearson ρ(embed cosine, distance)	-0.037	-0.137
Spearman ρ(embed cosine, distance)	-0.032	-0.150
Cramér's V (2×2 contingency)	0.032	0.083
Panel B: Quadrant Shares (%)		
High product × close geography	25.8	8.0
High product × far geography	24.2	5.9
Low product × close geography	24.2	39.4
Low product × far geography	25.8	46.7
Panel C: Conditional Probabilities		
P(close geo | high product)	0.516	0.576
P(close geo | low product)	0.484	0.457
Difference	0.032	0.119
P(high product | close geo)	0.516	0.169
P(high product | far geo)	0.484	0.112
Difference	0.032	0.057
Observations	31,939	184,160

Notes.—"High product proximity" ≡ embedding cosine similarity ≥ 0.609 (PE sample median). "Close geography" ≡ distance ≤ 998 km (PE sample median). Correlations computed at deal level between continuous embedding cosine and distance in km. Cramér's V measures the strength of association in the 2×2 contingency table (0 = independence). Sample restricted to deals with both company descriptions and valid ZIP codes. PitchBook add-on and corporate acquisition deals, 1990–2024.

Table 12: Baseline Rate-Channel Regressions

	Rate Level		Spread Level		Rate + Spread	
	(1)	(2)	(3)	(4)	(5)	(6)
PE × 10Y Rate	−0.1215***	−0.0620***			−0.1752***	−0.1160***
	(0.0064)	(0.0063)			(0.0071)	(0.0076)
PE × BAA-10Y Spread			−0.1176***	−0.0804***	−0.1738***	−0.1275***
			(0.0054)	(0.0059)	(0.0062)	(0.0072)
PE-Backed	0.4745***		0.6336***		1.3059***	
	(0.0241)		(0.0288)		(0.0458)	
N	496400	496165	496400	496165	496400	496165
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry FE	Yes	No	Yes	No	Yes	No
Acquirer FE	No	Yes	No	Yes	No	Yes
Cluster (acquirer)	Yes	Yes	Yes	Yes	Yes	Yes

Note: The dependent variable is the number of acquisitions by acquirer-year. The sample pools PE-backed platforms and non-PE corporate acquirers in the PitchBook acquirer-year rate-channel sample. The time fixed effect is calendar year, not quarter. Variables interacted with PE are standardized to unit variance. Odd columns include industry and year fixed effects; even columns include acquirer and year fixed effects. Standard errors are clustered by acquirer.

Table 13: Acquirer-Year Rate and Credit Sensitivity

Condition	Add-on freq. resp.	N
10Y yield	-0.0829***	496,165
	(0.0084)	
BAA–10Y spread	-0.0884***	496,165
	(0.0065)	
2Y yield	0.0142*	496,165
	(0.0076)	
EBP level	-0.0760***	496,165
	(0.0052)	
HY OAS	-0.0819***	495,727
	(0.0057)	
MPS	0.0677***	428,276
	(0.0053)	
Orth. MPS	0.0989***	428,276
	(0.0066)	
10Y–2Y term spread	-0.1123***	496,165
	(0.0076)	

Note: The dependent variable is the number of acquisitions by acquirer-year. Each row estimates equation (1) with acquirer and year fixed effects and standard errors clustered by acquirer. Financial condition variables are standardized using equal year weights. The Add-on frequency response column reports the raw PE-by-condition coefficient for a one-standard-deviation increase in the listed row variable. The EBP level row is reported separately from credit-spread levels. The 10Y–2Y term-spread row is the raw 10-year-minus-2-year spread, not a low-term-spread transformation. The MPS rows report raw positive short-rate surprises and are retained as shock diagnostics rather than financing-cost evidence. Sample years: 10Y yield, BAA–10Y spread, 2Y yield, EBP level, and 10Y–2Y term spread use 1995–2025; HY OAS uses 1996–2025; MPS and Orth. MPS use 1995–2023.

Table 14: PE Activity and National Industry Concentration

	$\Delta C^{(4)}$	$\Delta C^{(8)}$	$\Delta C^{(20)}$	$\Delta C^{(50)}$	ΔHHI
Panel A: PE add-on acquisitions, level count					
Add-ons$_{i,t-5\to t}$	-0.002063	-0.001276	-0.0002356	0.001365	-0.1201***
	(-1.36)	(-0.69)	(-0.10)	(0.47)	(-3.00)
Lagged concentration	-0.02394**	-0.009784	-0.01534	-0.0169*	-0.1667***
	(-2.00)	(-1.10)	(-1.52)	(-1.81)	(-3.42)
Panel B: buy-and-build platform entries, level count					
Entries$_{i,t-5\to t}$	-0.02679*	-0.02617	-0.02559	-0.01997	-0.9796***
	(-1.91)	(-1.50)	(-1.17)	(-0.75)	(-3.03)
Lagged concentration	-0.02689**	-0.01342	-0.01979*	-0.02212**	-0.1676***
	(-2.22)	(-1.49)	(-1.99)	(-2.46)	(-3.44)
Observations	1,322	1,322	1,322	1,322	814
Period fixed effects	Yes	Yes	Yes	Yes	Yes
Weights	Lag sales	Lag sales	Lag sales	Lag sales	Lag sales
SE clustering	NAICS-3	NAICS-3	NAICS-3	NAICS-3	NAICS-3

Notes: This table estimates a first-difference national concentration specification. Census concentration waves are 2012, 2017, 2022; five-year intervals are 2012–2017, 2017–2022. The regression is $\Delta C_{i,t}^{(n)} = \beta T_{i,t-5\to t} + \gamma C_{i,t-5}^{(n)} + \alpha_t + \epsilon_{i,t}$, using six-digit industries harmonized to strict one-to-one 2017 NAICS codes, lagged sales weights, and NAICS-3 clustered standard errors. Only level treatment counts are reported; log-one-plus specifications are excluded. The 2012 API table does not publish HHI, so the HHI column uses the 2017–2022 interval only. t statistics are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 15: Serial-Acquirer Exposure in Non-Empty DMI County-SIC-Year Markets

Presence category	Market-years	Share of non-empty markets
PE serial-platform exposure	20,191	0.0747%
Non-PE serial-acquirer exposure	29,636	0.1097%
Both PE and non-PE serial exposure	1,841	0.0068%
PE only	18,350	0.0679%
Non-PE serial only	27,795	0.1028%
Neither	26,978,568	99.8224%

Note: Denominator is non-empty DMI county-SIC-year markets with positive sales and employment. PE serial-platform exposure is computed from the DMI PE exposure panel and requires lagged platform roll-up add-ons of at least two. Non-PE serial-acquirer exposure is computed from PitchBook non-PE acquirer rows with at least two Merger/Acquisition deals, mapped to DMI county-SIC markets. The audit uses the full available DMI update through 2025; the main regression horizon is 1990–2024.

Table 16: PE Buy-and-Build Exposure and Local Market Concentration

	$\Delta C^{(4)}$	$\Delta C^{(8)}$	$\Delta C^{(20)}$	$\Delta C^{(50)}$
Panel A: any PE B&B HQ exposure				
Any PE B&B HQ exposure	0.688***	0.400***	0.168***	0.030
	(7.68)	(6.13)	(4.18)	(1.11)
Lagged concentration	-0.251***	-0.194***	-0.139***	-0.109***
	(-728.75)	(-499.84)	(-289.14)	(-170.87)
Panel B: platform-HQ exposure				
Platform HQ exposure	0.900***	0.512***	0.214***	0.045
	(7.61)	(6.12)	(4.10)	(1.25)
Lagged concentration	-0.251***	-0.194***	-0.139***	-0.109***
	(-728.71)	(-499.83)	(-289.14)	(-170.87)
Panel C: add-on-target HQ exposure				
Add-on HQ exposure	1.174***	0.793***	0.382***	0.109**
	(4.99)	(4.92)	(4.50)	(2.09)
Lagged concentration	-0.251***	-0.194***	-0.139***	-0.109***
	(-728.72)	(-499.83)	(-289.14)	(-170.87)
Observations	28,723,791	28,723,791	28,723,791	28,723,791
Market clusters	1,386,329	1,386,329	1,386,329	1,386,329
Within R-squared	0.111	0.085	0.059	0.044
Market fixed effects	Yes	Yes	Yes	Yes
State $\times$ year fixed effects	Yes	Yes	Yes	Yes
SIC-2 $\times$ year fixed effects	Yes	Yes	Yes	Yes
Market-size and acquisition controls	Yes	Yes	Yes	Yes
Regression weights	None	None	None	None
Period	1990–2024	1990–2024	1990–2024	1990–2024

Notes: This table estimates pure PE buy-and-build exposure local concentration regressions with no financial-condition interactions. The dependent variables are annual changes in firm-level sales-share concentration ratios in county-SIC-4 markets, in percentage points. Each panel estimates a separate exposure definition: any PE buy-and-build HQ exposure, lagged platform-HQ exposure, or add-on-target HQ exposure. All regressions include county-SIC market fixed effects, state-by-year fixed effects, SIC-2-by-year fixed effects, the outcome-specific lagged concentration ratio, log lagged establishments, log lagged employment, log lagged sales, lagged PE platform counts, PE add-on deal counts, non-PE M&A deal counts, and non-PE M&A exposure. Standard errors are clustered by county-SIC market. t statistics are in parentheses. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 17: Low-10Y Local Sales-Share CR4 Interaction Regressions

	(1) Any PE B&B HQ	(2) Platform HQ	(3) Add-on HQ
PE B&B HQ exposure	0.088	0.208	1.115***
	(0.65)	(1.30)	(3.62)
PE B&B HQ exposure x low 10Y	0.628***	0.754***	0.051
	(5.83)	(6.21)	(0.24)
Controls			
Control: Non-PE M&A exp. x low 10Y	0.474***	0.479***	0.514***
	(5.61)	(5.68)	(6.11)
Lagged sales-share CR4	-0.251***	-0.251***	-0.251***
	(-728.86)	(-728.83)	(-728.78)
Log lagged establishments	-2.458***	-2.458***	-2.458***
	(-328.21)	(-328.19)	(-328.10)
Log lagged employment	0.208***	0.208***	0.208***
	(70.39)	(70.38)	(70.36)
Log lagged sales	0.043***	0.043***	0.043***
	(22.42)	(22.42)	(22.43)
Lagged PE platforms	0.218***	-0.000	0.672***
	(2.98)	(-0.00)	(9.59)
PE add-on deal count	0.131	0.586***	-0.320
	(1.30)	(6.42)	(-1.58)
Non-PE M&A deal count	0.013	-0.013	0.006
	(0.54)	(-0.50)	(0.20)
Non-PE M&A exposure	-0.058	-0.033	-0.087
	(-0.61)	(-0.34)	(-0.89)
Observations	28,723,791	28,723,791	28,723,791
Market clusters	1,386,329	1,386,329	1,386,329
Within R-squared	0.111	0.111	0.111
Market fixed effects	Yes	Yes	Yes
State x year fixed effects	Yes	Yes	Yes
SIC-2 x year fixed effects	Yes	Yes	Yes
Period	1990–2024	1990–2024	1990–2024

Notes: This table reports the full low-10Y financial-condition interaction regression used in the rate-sensitive CR4 section. The dependent variable is the annual change in firm-level sales-share CR4 in a county-SIC-4 market, in percentage points. Columns differ only in the PE buy-and-build HQ exposure indicator: any PE buy-and-build HQ exposure, lagged platform-HQ exposure, or add-on-target HQ exposure. All regressions are unweighted county-SIC-year market cell regressions. They include county-SIC market fixed effects, state-by-year fixed effects, SIC-2-by-year fixed effects, lagged concentration and market-size controls, PE and non-PE acquisition controls, and the non-PE exposure-by-low-10Y placebo interaction. t statistics are in parentheses. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 18: Any-PE Local CR4 Interactions Across Financial Conditions

Financial condition	PE main	PE × condition	Years
Low 10Y yield	0.088	0.628***	1990–2024
Low 5Y yield	0.587***	0.173***	1990–2024
Low 2Y yield	0.695***	-0.014	1990–2024
Low 10Y–2Y term spread	0.674***	0.345***	1990–2024
Low BAA–10Y spread	0.772***	0.355***	1990–2024
Low GZ spread	0.856***	0.410***	1990–2024
Low EBP	0.657***	0.268***	1990–2024
Low HY OAS spread	0.462***	0.506***	1996–2024
Orthogonalized MPS easing	0.630***	-0.373***	1990–2023

Note: Dependent variable is the annual change in firm-level sales-share CR4. Each row uses any-PE-HQ exposure and replaces the financial-condition interaction in the baseline specification. For yield, spread, and policy-shock variables, the condition is constructed as $-z(X_t)$, so larger values correspond to lower yields, lower spreads, or easier policy surprises. Treasury-yield rows use annual means of daily constant-maturity yields; the local 2Y and 5Y files were extended for 1990–1994 using the U.S. Treasury daily par-yield-curve archive, which also extends the 10Y–2Y term-spread row to the full 1990–2024 CR4 window. The HY row is already the ICE BofA U.S. High Yield Index option-adjusted spread (OAS), so no separate OAS row is added; it remains source-limited because the series begins at the end of 1996. The MPS row uses annual sums of Bauer–Swanson orthogonalized short-rate surprises and ends in 2023. Table 19 defines each proxy.

Table 19: Financial-Condition Proxy Definitions

Proxy	Construction	Source	Economic interpretation
Low 10Y yield	$-z(\overline{DGS10}_t)$, where $\overline{DGS10}_t$ is the annual mean 10-year Treasury yield	FRED DGS10 / U.S. Treasury constant-maturity yield	Lower long risk-free rate and lower discount-rate/debt-service environment
Low 5Y yield	$-z(\overline{DGS5}_t)$, annual mean 5-year Treasury yield; 1990–1994 backfilled from Treasury daily par-yield data	FRED DGS5 and U.S. Treasury daily par-yield curve	Medium-term risk-free financing conditions
Low 2Y yield	$-z(\overline{DGS2}_t)$, annual mean 2-year Treasury yield; 1990–1994 backfilled from Treasury daily par-yield data	FRED DGS2 and U.S. Treasury daily par-yield curve	Shorter-maturity risk-free rate, closer to expected near-term policy
Low 10Y–2Y term spread	$-z(\overline{DGS10}_t - \overline{DGS2}_t)$	FRED/Treasury 10Y and 2Y Treasury yields	Lower term spread or flatter yield curve, distinct from a low short-rate level
Low BAA–10Y spread	$-z(\overline{BAA10Y}_t)$	FRED BAA10Y, Moody's BAA corporate yield minus 10Y Treasury yield	Lower investment-grade credit spread and lower risky-debt financing premium
Low GZ spread	$-z(GZ_t)$	Gilchrist–Zakrajsek corporate-bond spread data (Gilchrist and Zakrajšek, 2012)	Lower default-risk credit-spread component and less distress/valuation pressure
Low EBP	$-z(EBP_t)$	Gilchrist–Zakrajsek excess bond premium (Gilchrist and Zakrajšek, 2012)	Lower residual credit premium, healthier intermediary balance sheets, and looser credit supply
Low HY OAS spread	$-z(HYOAS_t)$	FRED/ICE BofA U.S. High Yield Index option-adjusted spread (BAMLH0A0HYM2)	Lower leveraged-credit spread relevant for riskier buyout financing
Orthogonalized MPS easing	$-z(\sum_{m\in t} MPS_m^{orth})$	Bauer–Swanson monthly monetary-policy surprises (Bauer and Swanson, 2023)	High-frequency short-rate policy easing, purged of pre-FOMC macro/market information

Note: All variables are annual and standardized before interaction with any-PE-HQ exposure. The sign convention is chosen so a positive "condition" value consistently means easier financial conditions in Table 18.

Table 20: Local/Non-Tradable DMI CR4 Results

	(1) Full	(2) Local/NT
PE B&B	0.088	0.754***
	(0.65)	(2.81)
PE B&B × low 10Y	0.628***	0.649***
	(5.83)	(2.85)
Controls		
Non-PE M&A × low 10Y	0.474***	0.499***
	(5.61)	(3.03)
Lag CR4	-0.251***	-0.261***
	(-728.86)	(-593.11)
Log lag estabs.	-2.458***	-2.907***
	(-328.21)	(-265.00)
Log lag employ.	0.208***	0.239***
	(70.39)	(42.73)
Log lag sales	0.043***	0.110***
	(22.42)	(28.34)
Lag PE platforms	0.218***	0.263***
	(2.98)	(2.73)
PE add-on deals	0.131	0.012
	(1.30)	(0.08)
Non-PE M&A deals	0.013	-0.001
	(0.54)	(-0.05)
Non-PE M&A	-0.058	0.222
	(-0.61)	(1.17)
Observations	28,723,791	14,546,906
Market clusters	1,386,329	586,281
Within R-squared	0.111	0.114
Market FE	Yes	Yes
State × year FE	Yes	Yes
SIC-2 × year FE	Yes	Yes
Period	1990–2024	1990–2024

Note: Dependent variable is the annual change in firm-level sales-share CR4. Local/NT denotes SIC markets in the paper's local-service industry definition. PE B&B presence equals either platform-HQ or add-on-HQ presence in the county-SIC-year. All regressions are unweighted and include county-SIC market, state-by-year, and SIC-2-by-year fixed effects. t statistics are in parentheses. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 21: Mechanical PE Contributions to Local Sales-Share CR4

Sample	Exposure	Main pp/yr	Amp. pp/yr	Cum. total pp
All industries	Any PE B&B HQ	0.0058	0.0333	1.369
All industries	Platform HQ	0.0110	0.0366	1.667
All industries	Add-on HQ	0.0143	0.0006	0.524
Local/non-tradable	Any PE B&B HQ	0.0446	0.0365	2.841
Local/non-tradable	Platform HQ	0.0532	0.0418	3.325
Local/non-tradable	Add-on HQ	0.0352	-0.0023	1.152

Note: Current-sales-weighted mechanical accounting using the CR4 low-10Y regressions. "Main pp/yr" is the PE main component and "Amp. pp/yr" is the low-10Y interaction component. Cumulative totals sum annual contributions over 1990–2024. Weights are normalized across all county-SIC markets in each sample-year (all industries or local/non-tradable), including markets without PE exposure; the totals are not conditional effects within PE-exposed markets.

Table 22: Affected-Market Any-PE Counterfactual Diagnostic

Metric	All industries	Local/non-tradable
Affected markets (%)	0.116	0.093
Affected sales (%)	5.8	5.9
Main contribution (pp/yr)	0.1003	0.7545
Amplification contribution (pp/yr)	0.3345	0.3491
Total contribution (affected markets only, pp/yr)	0.4348	1.1036
Cumulative total (affected markets only, pp)	15.219	38.624

Note: Affected markets are county-SIC-year markets with any PE buy-and-build exposure. The first two columns report the affected set's average annual unweighted market share and current-sales-weighted share relative to the full sample-year denominator. Contribution rows, including the two affected-markets-only total rows, re-normalize current-sales weights only over affected markets. The diagnostic is conditional on affected markets and complements, rather than replaces, the sample-wide aggregate counterfactual in Table 21.

Table 23: Local Economic Outcome Magnitudes Across QCEW, SUSB, QWI, and BDS

Outcome	Granularity	Period	PE coefficient	P90 treated effect
QCEW Δ log establishments	county–N6–quarter	1990Q2–2024Q4	0.0001***	0.81%
QCEW Δ log employment	county–N6–quarter	1990Q2–2024Q4	-0.0000	-0.01%
QCEW Δ log wage per employee	county–N6–quarter	1990Q2–2024Q4	0.0001	0.49%
SUSB Δ log firms	state–N6–year	2009–2021	0.0003***	0.58%
SUSB Δ log establishments	state–N6–year	2009–2021	0.0000	0.02%
SUSB Δ log employment	state–N6–year	2009–2021	-0.0002	-0.46%
QWI Δ log employment	county–N4–year	2009–2024	0.0010	0.56%
QWI Δ log stable earnings	county–N4–year	2009–2024	0.0003	0.17%
QWI Δ log hires	county–N4–year	2009–2024	0.0013	0.66%
QWI Δ log separations	county–N4–year	2009–2024	0.0032**	1.64%
QWI Δ log turnover	county–N4–year	2009–2024	-0.0004	-0.17%
QWI Δ log firm job gains	county–N4–year	2009–2024	0.0022*	1.16%
QWI Δ log firm job losses	county–N4–year	2009–2024	0.0021**	1.12%
BDS Δ log firms	county–sector–year	1990–2023	0.0003*	0.21%
BDS Δ log establishments	county–sector–year	1990–2023	0.0000	0.03%
BDS Δ log employment	county–sector–year	1990–2023	-0.0006	-0.51%
BDS Δ log establishment entry	county–sector–year	1990–2023	0.0004	0.29%
BDS Δ log establishment exit	county–sector–year	1990–2023	-0.0018	-1.20%
BDS Δ log firm deaths	county–sector–year	1990–2023	-0.0021	-1.22%
BDS Δ log job creation	county–sector–year	1990–2023	-0.0004	-0.33%
BDS Δ log job destruction	county–sector–year	1990–2023	-0.0007	-0.58%

Note: Coefficients are baseline PE-add-on-intensity coefficients from source-specific fixed-effect regressions, expressed per disclosed PE add-on per 1,000 lagged denominator units in each panel. P90 treated effect evaluates the coefficient at the 90th percentile of positive PE intensity. The granularity column reports the geography–industry–frequency available in each source; N6 and N4 denote NAICS-6 and NAICS-4, respectively. The SUSB firm-count row corresponds to the state–NAICS-6 firm-count surface shown in Online Appendix Figure A.5. The public SUSB state–NAICS-6 and county–NAICS-3 files used in this project report stock variables–firms, establishments, employment, payroll, and, in the 2022 county file, receipts–but not flow variables such as firm entry, firm births, firm deaths, establishment entry, or establishment exit. Those entry/death margins are therefore available here only in the coarser BDS county–sector panel. Periods report the best feasible span inside the 1990–2024 target horizon: QCEW reports the full available 1990Q2–2024Q4 private panel; SUSB and QWI begin when public-panel coverage and first differences are available; and BDS is available through 2023. The 2001Q1–2024Q4 support-stable QCEW sensitivity in Online Appendix A.11, motivated by the thinner county–NAICS-6 support in BLS reconstructed 1990–2000 rows, yields very similar baseline P90 effects: 0.78% for establishments, -0.01% for employment, and 0.58% for wage per employee. Public QWI does not provide an establishment-count outcome in this panel; SUSB and BDS stock outcomes are included where public granularity permits.

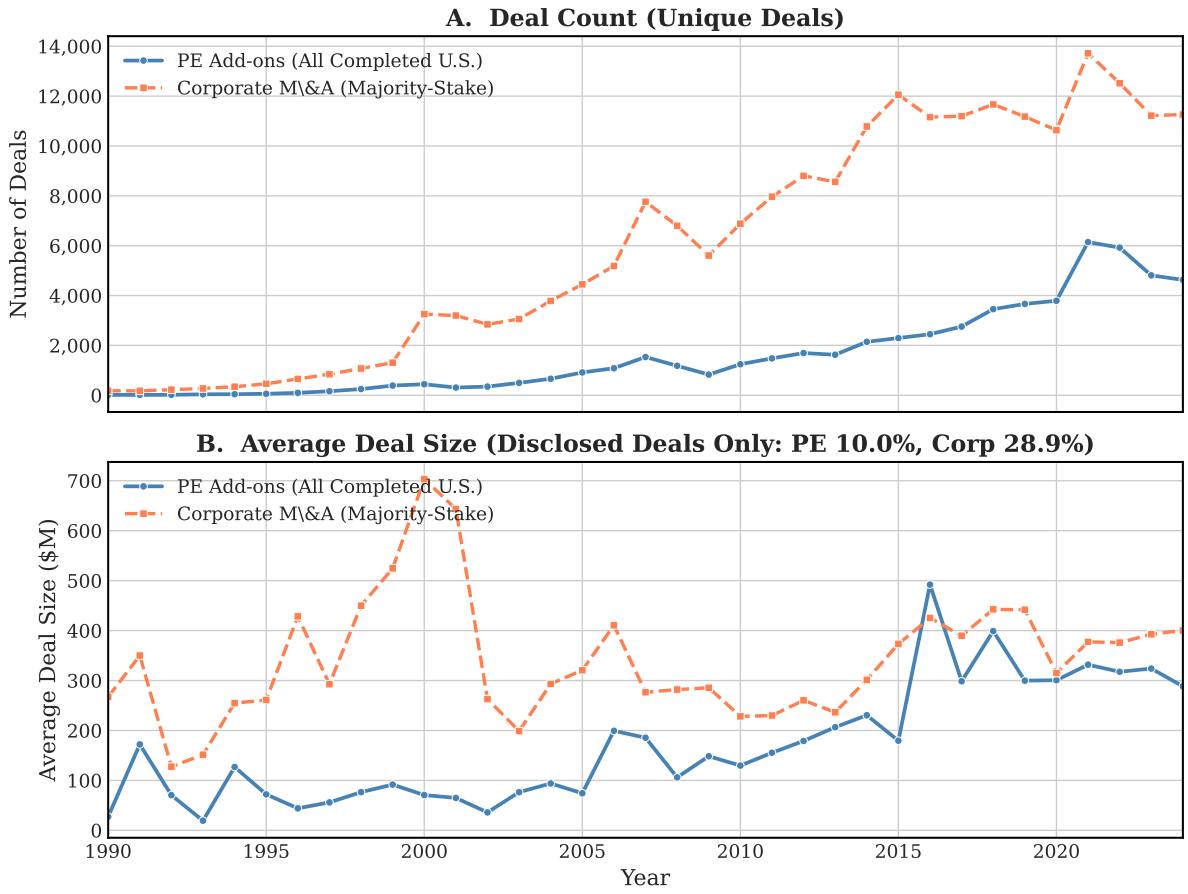


Figure 1: PE Add-On Acquisitions and Corporate M&A: Deal Count and Average Size

Note: The PE add-on series uses all completed U.S. add-on deals in PitchBook's North America Private Equity universe, before the platform-link and holding-window filters used for the main analysis sample. The comparison series uses non-PE majority-stake corporate control acquisitions, 1990–2024. The non-PE control benchmark keeps Merger/Acquisition, Corporate Asset Purchase, and Buyout/LBO deal types with majority control acquired or undisclosed control percentages, and excludes PE-backed transactions. Panel A reports unique deal counts by year. Panel B reports average deal size among transactions with disclosed deal values; deal-value coverage is 10.0 percent for PE add-ons and 28.9 percent for non-PE control acquisitions.

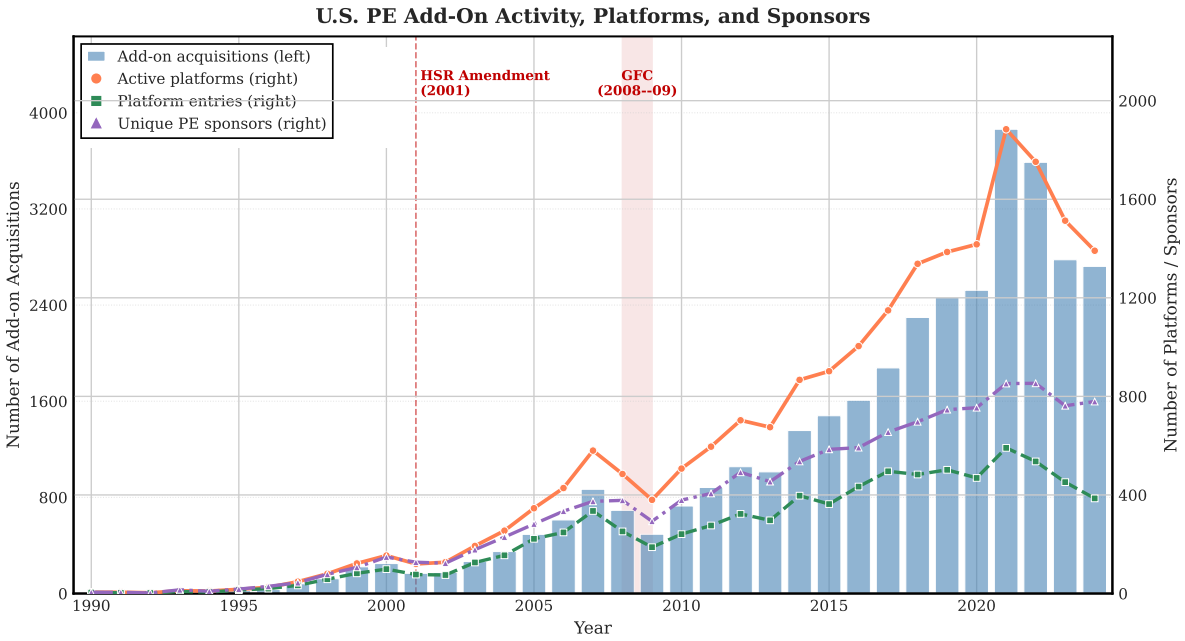
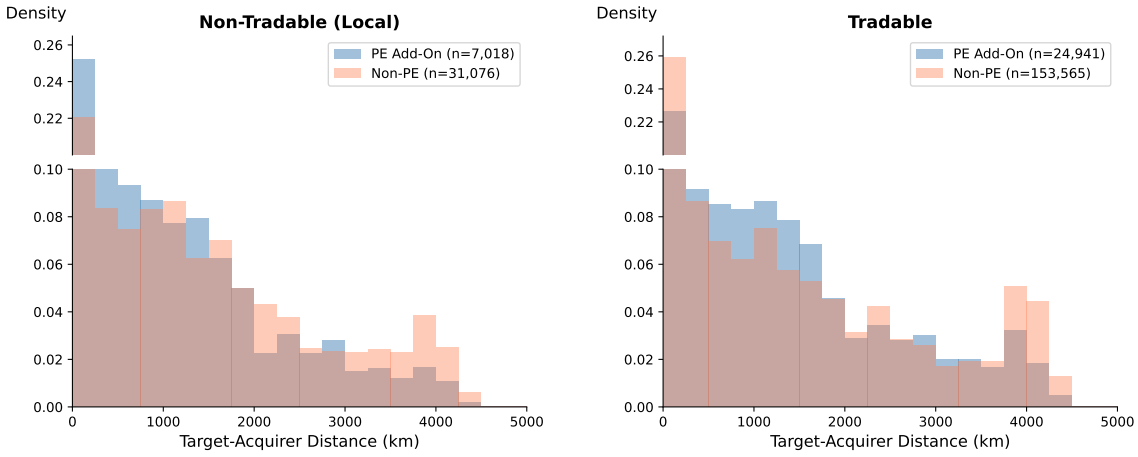


Figure 2: U.S. PE Add-On Activity, 1990–2024

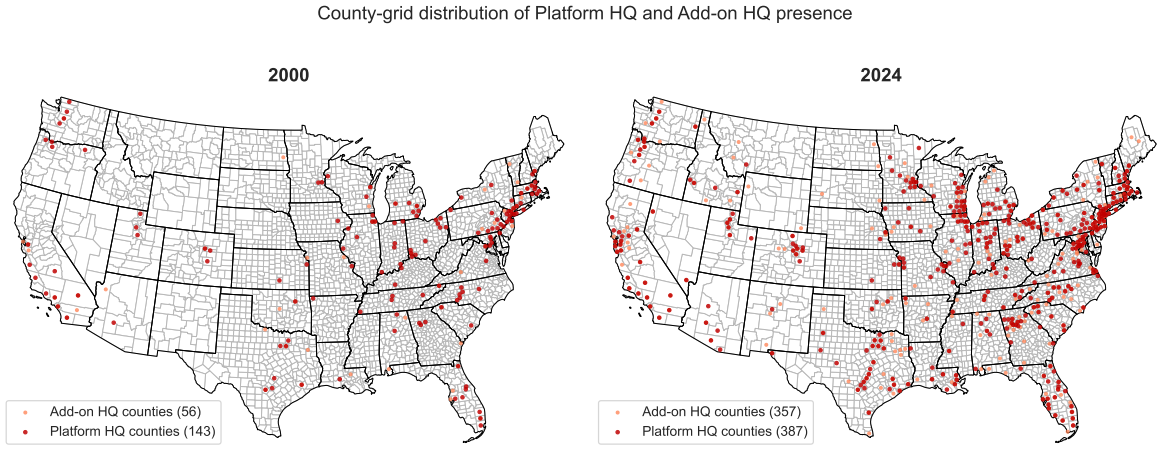
Note: Bars report annual completed U.S. PE add-on acquisitions in the analysis-window sample. Lines report platforms with at least one add-on in the year, platform entries measured by the first observed add-on year, and unique PE sponsors named on add-on transactions. Sample: 35,118 add-ons linked to 8,256 PE-backed platform companies from the PitchBook North America Private Equity universe. The dashed vertical line marks the 2001 Hart–Scott–Rodino notification-threshold revision; the shaded band marks the 2008–09 global financial crisis.

Figure 3: Target–Acquirer Distance by Industry Tradability



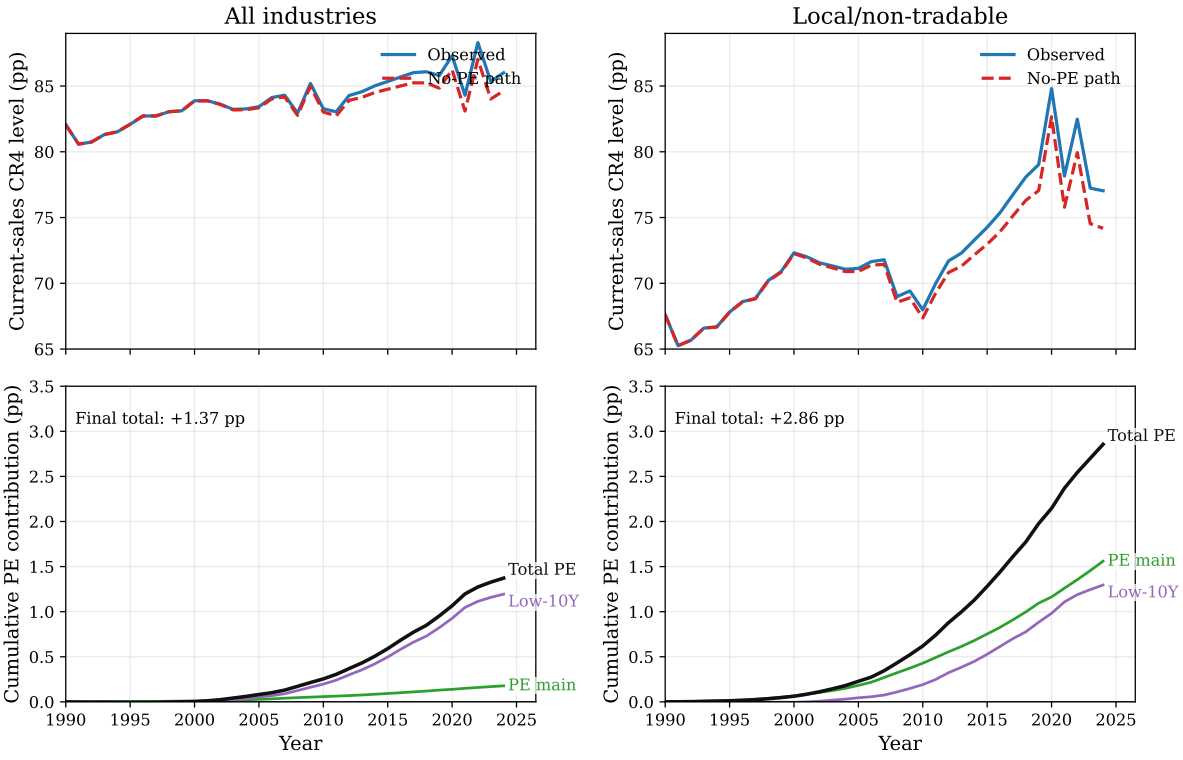
Note: Kernel density of haversine distance between target and acquirer ZIP-5 centroids, split by target-industry tradability. The non-tradable split is the relevant visual diagnostic for local-market competition.

Figure 4: County-Grid Distribution of Platform-HQ and Add-on-HQ Exposure



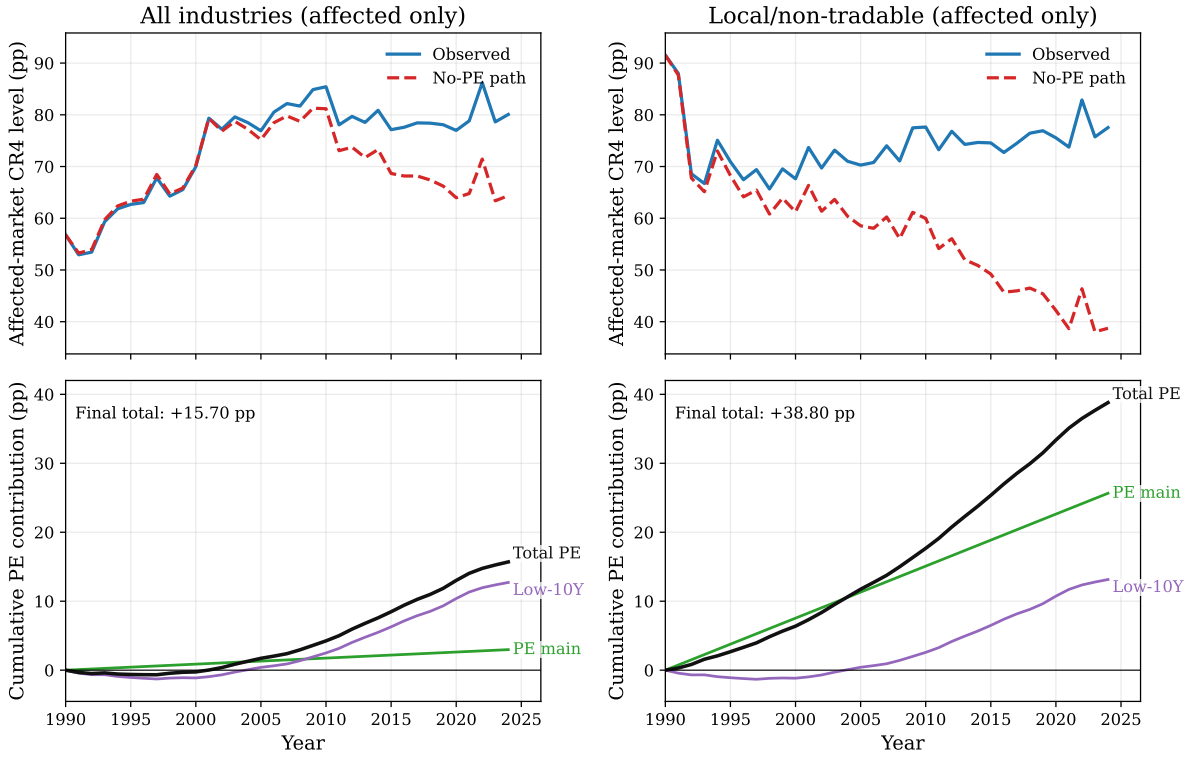
Note: County boundaries use Census county polygons for the contiguous U.S. Dots are county centroids for counties with at least one exposed county-SIC market in the indicated year. Platform-HQ exposure is shown in red and add-on-HQ exposure in orange.

Figure 5: Time-Series Impact of Any-PE Counterfactual Contributions



Note: Top panels plot the observed current-sales-weighted CR4 level $C_t^{obs} = \sum_m w_{m,t} C_{m,t}$ and the mechanical no-PE path $C_t^{noPE} = C_t^{obs} - \sum_{\tau=t_0+1}^{t} \Delta C_\tau^{PE}$, anchored to the observed level in the first plotted year t_0. The annual PE contribution is $\Delta C_t^{PE} = \sum_m w_{m,t}[\widehat{\beta} PE_{m,t} + \widehat{\theta}(PE_{m,t} \times Low10Y_t)]$, where $w_{m,t}$ is the market's current-sales share within the all-industry or local/non-tradable sample-year, $PE_{m,t}$ is any PE buy-and-build HQ exposure, and $Low10Y_t$ is the standardized low-10-year-yield condition. Bottom panels cumulate the $\widehat{\beta} PE_{m,t}$ main component, the $\widehat{\theta}(PE_{m,t} \times Low10Y_t)$ amplification component, and their total. Untreated markets receive zero mechanical PE contribution; the exercise does not model entry, prices, quality, ownership restructuring, or general-equilibrium responses.

Figure 6: Affected-Market Time-Series Impact of Any-PE Counterfactual Contributions



Note: This figure repeats Figure 5 after restricting each sample-year to affected county-SIC-year markets, defined as markets with any PE buy-and-build HQ exposure. Current-sales weights are re-normalized within that affected set, so the top panels plot affected-market CR4 levels and affected-market no-PE paths rather than sample-wide aggregate paths. The same annual contribution equation is used, but the summation is over affected markets only. Bottom panels cumulate the affected-market PE main component, low-10Y amplification component, and total contribution.