Datacenter Mortgages and the Originate-to-Distribute Model*

Pedro Gete†, Amparo Mercader‡

August 2026

Abstract

We build a novel property-level database of U.S. datacenter mortgages, an opaque AI-buildout financing channel. Linking mortgages to lender balance sheets, we document a large, concentrated market dominated by originate-to-distribute (OTD) banks and Undiversified Datacenter borrowers, whose repayment depends on continued AI demand rather than diversified revenue. Across complementary identification strategies, we find that the High-OTD leverage gap widens more than fourfold and that High-OTD lenders disproportionately finance these borrowers. Moreover, the leverage premium itself concentrates among Undiversified Datacenter borrowers, linking distribution incentives to looser underwriting and greater exposure to AI-demand risk.

Keywords: Artificial Intelligence; Datacenters; Leverage; Securitization; Originate-to-Distribute; Digital Infrastructure.
JEL Classification: G12, G21, G23, R31.

*We thank Olivier Darmouni, Carlos Garriga, Lucas Gulão, Carlos Ramirez, Victoria Vanasco, Stijn Van Nieuwerburgh, Isabel de Valenzuela, and participants at the 2026 Dallas Fed Conference, the University of Alicante, and the 2026 Spanish Finance Conference for helpful comments. This research was partially supported by MICIU/AEI/10.13039/501100011033/FEDER, UE, Grant No. PID2024-161318NB-I00. The authors have no financial relationships or other potential conflicts of interest to disclose.

†IE University. Email: pedro.gete@ie.edu.
‡IE University and Georgetown University. Email: mercada@georgetown.edu

1 Introduction

The AI-driven datacenter buildout draws on a wide range of financing channels, including corporate bonds and equity, private credit, vendor financing, and property-level mortgages secured by datacenter assets. Existing work focuses primarily on corporate- and sponsor-level financing channels (Van Nieuwerburgh, 2026). In this paper we study property-level datacenter mortgages, a largely opaque segment. These loans are fragmented across property records, local lenders, and non-public loan documents, with no centralized disclosure comparable to residential mortgage or corporate bond markets. As a result, little is known about the size of this market, the institutions financing it, or the lending standards applied to these loans.

We construct a novel property-level dataset linking U.S. datacenter properties and mortgage originations to lender balance sheets. The resulting data allow us to observe property characteristics, mortgage terms, and lender business models jointly and, therefore, to study underwriting standards through both leverage and borrower selection. To our knowledge, this is the first property-level dataset linking U.S. datacenter mortgages to lender balance sheets.

The data reveal a large, recent, and highly concentrated market. Datacenter mortgage originations were negligible before 2021 but accelerated abruptly in mid-2021. By May 2026, the sample contains approximately $138 billion in allocated mortgage debt, roughly 84% of which is financed by distribution-oriented banks, and the lending boom is directed primarily toward Undiversified Datacenter borrowers, that is, firms whose repayment capacity depends on continued demand for datacenter capacity rather than diversified sources of revenue.

This concentration is economically important because the demand supporting these borrowers may be transitional rather than durable. Large technology firms have described their reliance on third-party datacenter capacity as a bridge while internal capacity comes online.[1] A reversal in demand for leased capacity would therefore fall disproportionately on

[1] Alphabet's CFO described external capacity as a stopgap to be scaled back as internal capacity comes

Undiversified Datacenter borrowers, which lack diversified sources of repayment. Their exposure makes this market a natural setting in which to study whether lender incentives affect underwriting standards.

The originate-to-distribute (OTD) literature provides a natural framework for this question. The 2008 financial crisis showed that lenders expecting to sell or securitize loans had weaker incentives to screen borrowers and originated riskier mortgages (Keys et al., 2010; Purnanandam, 2011); related concerns extend to syndicated corporate credit (Bord and Santos, 2015). We ask whether these incentives reappear in a new asset class created by the AI-driven investment boom. Following Purnanandam (2011), we construct bank-level OTD scores using the trading-asset and repo measures of market-based intermediation proposed by Adrian and Shin (2010), and classify banks in the top two deciles as High-OTD.

Our empirical strategy exploits the abrupt acceleration in datacenter lending beginning in mid-2021, with July 2021 marking the empirical onset of the lending boom. In a difference-in-differences (DiD) design, we compare changes in underwriting outcomes at properties financed by High-OTD banks with those financed by other lenders before and after the onset of the lending boom.

We find that distribution-oriented lenders substantially relax lending standards as the datacenter market expands. Following the onset of the lending boom, the High-OTD leverage gap widens by more than fourfold across measures of debt per square foot, debt-to-rent, and loan-to-value. High-OTD lenders also disproportionately finance Undiversified Datacenter borrowers, and the leverage premium is concentrated among these same borrowers. Distribution-oriented lenders thus respond to the expansion of lending opportunities along both margins of underwriting. They extend more leverage and channel credit toward borrowers most vulnerable to shifts in datacenter demand.

We support the baseline DiD identification with a set of tests targeting its leading alternative explanations. Event-study estimates show no differential pre-trends before the lending

online. Alphabet Q2 2026 Earnings Call, July 22, 2026.

boom. Matching and borrower fixed-effects specifications address sorting across borrowers and projects, while estimates from a non-datacenter commercial real estate placebo reverse sign, ruling out generic lending aggressiveness by High-OTD banks. Results are also robust to predetermined bank-size controls, exclusion of the largest lenders, alternative OTD classifications, and alternative constructions of leverage. We then complement the DiD with an instrumental-variables strategy based on predetermined bank balance sheets, which yields a strong first stage and a null pre-boom placebo. Taken together, these complementary approaches provide consistent evidence that OTD incentives shape underwriting behavior.

The remainder of the paper proceeds as follows. Section 2 reviews the related literature. Section 3 describes the data, the construction of the OTD classification, and our measures of lending standards. Section 4 documents the datacenter mortgage market. Section 5 presents the baseline DiD analysis. Section 6 evaluates alternative explanations and robustness. Section 7 develops the instrumental-variables strategy; Section 8 studies the loan-size mechanism; and Section 9 concludes.

2 Related Literature

This paper connects two rapidly converging literatures. The first studies the OTD model and how distribution incentives affect lending standards. The second examines the economics of the AI-driven datacenter buildout. Our contribution is to build a novel property-level dataset to ask whether the distribution incentives that proved central during the 2008 financial crisis re-emerge in a new asset class created by the AI boom.

The OTD literature argues that lenders expecting to distribute rather than retain credit have weaker incentives to screen and monitor borrowers because loan quality has a smaller effect on their future income (Akerlof, 1970; DeMarzo and Duffie, 1999; Gorton and Pennacchi, 1995; Stiglitz and Weiss, 1981). The 2008 financial crisis brought this mechanism to the forefront. Residential mortgage studies link securitization to weaker screening, higher de-

fault risk, and looser credit supply (Berndt and Gupta, 2009; Dell'Ariccia et al., 2012; Keys et al., 2010; Loutskina, 2011; Loutskina and Strahan, 2009; McGowan and Nguyen, 2023; McGowan et al., 2024; Mian and Sufi, 2009; Piskorski et al., 2010; Purnanandam, 2011). Related evidence from corporate lending shows that distribution-oriented arrangers impose weaker covenants, select riskier borrowers, and monitor less intensively (Bord and Santos, 2015; Gustafson et al., 2021; Ivashina and Sun, 2011; Puri et al., 2011; Sufi, 2007), while commercial real estate studies document similar patterns for CMBS conduit lending (Black et al., 2012). Although post-crisis regulation reduced some of these incentives, commercial datacenter mortgages remain largely outside the regulatory framework designed for residential securitization, and evidence suggests that OTD incentives continue to shape commercial syndicated lending (Acharya et al., 2013; Aobdia et al., 2021; Bord and Santos, 2015).

At the same time, a rapidly expanding literature studies the economics of the AI-driven datacenter boom. Existing work examines the financing architecture of the AI buildout (Van Nieuwerburgh, 2026), state tax incentives for datacenter investment (Gargano and Giacoletti, 2025), the response of renewable-energy investment to datacenter electricity demand (Darmouni et al., 2025; Heiss et al., 2026), and the local economic effects of datacenter entry (Yue and Zeng, 2026). Collectively, these papers explain where datacenters are built, how they are financed, and how they affect energy systems and local economies.

3 Database

3.1 Data Sources

We construct a novel property-loan-level dataset of U.S. datacenters by integrating commercial real estate records, loan origination data, and lender balance sheet information. The dataset links leveraged datacenter properties to their mortgage lenders, enabling us to measure underwriting outcomes at the property level.

The core of the dataset comes from CoStar, which provides a national inventory of U.S.

datacenter properties as of May 2026 (2,719 assets). From this universe, we identify 668 properties with observed mortgage debt (25% of the universe), providing broad national coverage of the financed datacenter market. These properties form the mortgaged sample used in the leverage analysis. Many remaining properties are owned by large technology firms that can finance investment through corporate debt rather than property-level mortgages. Datacenter capacity is mapped across U.S. counties in the appendix (Appendix Figure A1).

For each property, we observe rentable building area (RBA), geographic location, development status, shell rent, property taxes, percent leased, year built, borrower information, and loan origination records. CoStar's loan records include origination amount, origination date, originating lender name, and other loan characteristics.

We link these properties to their lenders' quarterly Call Report data, which provides the balance-sheet information underlying our OTD classification. We further complement the mortgage data with DealScan records of syndicated facilities from 2015 through March 2026. Unlike CoStar's property-level mortgage data, DealScan includes only syndicated facilities but covers a broader range of debt types, not just mortgages; we use it to validate the OTD classification against an independent measure of lender behavior.

Our baseline empirical analysis uses the full sample of 668 mortgaged datacenter properties; some specifications rely on different samples because of data availability or empirical design, as detailed where each arises. Appendix A.1 details data sources and sample construction, and Appendix A.2 describes how the estimation sample varies across specifications.

3.2 OTD Classification

We classify lenders according to their propensity to originate loans for distribution rather than retention. Following the OTD framework of Purnanandam (2011), we construct a balance-sheet measure of distribution orientation using the market-based intermediation variables proposed by Adrian and Shin (2010). Banks receive higher OTD scores when they hold a smaller share of assets as loans and larger shares as trading assets and repurchase

agreements, reflecting a business model oriented toward originating and distributing credit rather than holding it on balance sheet. We rank banks with Call Report scores according to their average score over 2017–2020 and classify those in the top two deciles as *High-OTD*. Appendix A.3 describes the classification of other lenders. This window is predetermined relative to the July 2021 shock that anchors our identification strategy, to account for the possibility that a bank's business model could shift because of the datacenter lending boom rather than only cause the leverage response to it; Section 6 shows the results are robust to an alternative classification window.

We validate the classification using DealScan, an independent dataset of syndicated credit facilities restricted to domestic banks. Domestic High-OTD banks average $636.5 billion in matched DealScan volume per bank, against $88.7 billion for Low-OTD banks, and are less likely to participate via full syndication and more likely to participate in club deals[2] (Appendix Table B1), a partial out-of-sample confirmation of the OTD label.

3.3 Measuring Lending Standards

Lending standards are multidimensional, so this section constructs measures for each of the two dimensions studied throughout the paper, beginning with leverage.

Leverage measures. We construct property-level leverage measures from loan origination records. When a single loan is secured by multiple properties, we allocate the loan balance across collateral properties in proportion to each property's RBA. This allows us to measure debt per square foot and debt-to-rent ratios at the property level in a manner comparable across single-asset and portfolio loans. Table 1 reports summary statistics for the main variables. Appendix A provides complete variable definitions and data cleaning procedures.

Unlike traditional commercial real estate, where mortgage debt is used to finance land and building structure, datacenter loans are also used to finance specialized equipment,

[2]Full syndication means the loan is distributed broadly across many participant lenders. A club deal is shared among a small group of relationship lenders without wide distribution. The lead arranger is the bank that structures and coordinates the syndicate.

GPUs, cooling, and network infrastructure. Because specialized equipment is not separately observed or verified at the property level in our data, we measure collateral using real estate value alone throughout the paper, excluding the technology layered on top of it. This valuation approximates what a lender would recover if AI demand for datacenter capacity collapsed and the specialized equipment were stranded, leaving only the underlying real estate to service any remaining loan balance over the life of the loan.

To measure real-estate collateral values empirically, we construct three valuation approaches for each property: government tax assessments from matched parcel records, discounted cash flow values capitalizing shell rent, and acre-based land values using an implied dollars-per-acre measure from same-state assessments. Loan-to-value ratio is defined as mortgage debt divided by real-estate collateral value, taken as the highest of the three approaches, reflecting the property's highest-value use.

Borrower classification. We turn next to the second dimension of lending standards, borrower selection. This dimension captures which borrowers receive financing from High-OTD lenders. We classify borrowers according to the diversification of their revenue base. *Undiversified Datacenter* borrowers are firms whose primary business is owning and operating datacenter real estate. Because their revenues are concentrated in datacenter leasing and related services, their repayment capacity depends directly on AI-driven demand for datacenter capacity, with no diversified income to fall back on should that demand recede. This is the borrower group that our borrower-selection analysis follows throughout the paper, precisely because its revenue concentration makes it the group most exposed to a pullback in AI demand. The remaining borrowers are classified as *Big Tech*, *Government*, or *Other*, which includes REITs, institutional investors, and other diversified corporations. Classification is based on each entity's known primary line of business: firms whose disclosed operations center on owning and operating datacenter real estate are coded Undiversified Datacenter, and all others are assigned by entity type. Table 2, Panel A, reports summary statistics by borrower type; the full list of named borrowers in each category is in Appendix Table B2.

4 The Datacenter Mortgage Market

We begin by describing the new datacenter mortgage database. The data reveal three defining characteristics of this market. First, property-level datacenter mortgage lending is a recent phenomenon, expanding rapidly after mid-2021. Second, this expansion is highly concentrated among distribution-oriented lenders. Third, mortgage financing is concentrated primarily among Undiversified Datacenter borrowers. These descriptive patterns motivate the empirical analysis that follows.

4.1 Market Growth and Lender Composition

Figure 1 documents the evolution of datacenter mortgage originations by lender type. Through mid-2021 all lender types originate at modest rates, after which mortgage origination accelerates abruptly. We use July 2021, when this acceleration begins, as the empirical onset of the lending boom. Distribution-oriented (High-OTD) banks account for the dominant share of this expansion, while Low-OTD banks, non-bank lenders, and government lenders expand much more gradually. The rapid rise of High-OTD lending establishes the setting in which OTD incentives are most likely to operate.

4.2 Borrower Composition

Panel B of Figure 1 shows that the lending boom is concentrated primarily among Undiversified Datacenter borrowers, firms whose primary business is datacenter leasing. Unlike diversified corporations, these borrowers depend almost exclusively on continued AI-driven demand for datacenter capacity to service their debt. Many datacenter properties owned by large technology firms are financed through other channels, such as unsecured corporate debt, and therefore fall outside our property-level mortgage sample.[3] Consequently, our

[3]Big Tech firms own 26% of total datacenter space in our sample yet account for just 0.2% of mortgage debt outstanding, since they finance primarily on their own corporate balance sheets, through bonds and equity, rather than property-level mortgages.

database is not intended to represent the universe of datacenter owners, but rather the segment financed through property-level collateralized lending. This distinction is economically important because repayment depends directly on the value and utilization of the underlying datacenter assets, making collateral quality central to credit risk and to lenders' underwriting decisions. Undiversified Datacenter borrowers carry approximately 2.3 times more debt per square foot than all other borrowers.

Together, these descriptive patterns identify a rapidly expanding mortgage market dominated by distribution-oriented lenders and Undiversified Datacenter borrowers. The remainder of the paper asks whether the concentration of lending among distribution-oriented institutions is associated with systematically weaker underwriting standards.

5 Difference-in-Differences

Our baseline empirical strategy is a DiD specification that exploits the abrupt acceleration in datacenter lending beginning in July 2021. This design invites several natural objections, including sorting across borrowers or projects, generic lending aggressiveness at High-OTD banks, bank size confounds, and the endogeneity of OTD status itself. Section 6 maps each of these threats to the specific tests that address it.

5.1 Identification Strategy

The treatment group consists of properties financed by High-OTD banks, while the control group consists of properties financed by all other lenders. The identifying assumption is that, absent the AI-driven datacenter lending boom, underwriting outcomes at properties financed by High-OTD and other lenders would have followed parallel trends.

We evaluate whether distribution-oriented lenders relax lending standards, measured through leverage and borrower selection, following the AI-driven datacenter lending boom. We begin by estimating an event-study specification to assess the parallel-trends assumption

underlying the DiD specification. We then estimate the corresponding DiD specification to quantify the average post-shock effect on lending standards.

The event-study specification is

$$Y_{it} = \beta_H \, \text{HighOTD}_i + \sum_{p \neq -1} \beta_p \left(\text{HighOTD}_i \times \mathbf{1}\{t \in p\} \right) + \gamma_s + \delta_t + \eta_p + \varepsilon_{it}, \tag{1}$$

where p indexes these periods, with period -1 omitted as the reference category in estimation. State fixed effects are denoted by γ_s, loan-origination-year fixed effects by δ_t, and 15-month period fixed effects by η_p; because each period spans parts of two calendar years, the two sets of time fixed effects are not collinear. Standard errors are clustered at the lender level. After estimation, we re-center the reported coefficients relative to the average High-OTD gap over the two periods immediately before the shock (-1 and -2), and estimate the specification for three leverage outcomes, ln(DebtPSF), ln(DebtToRent), and ln(LTV).

Event studies of lending standards and the loan-size mechanism include state, origination-year, and event-window fixed effects. The borrower-selection event study is estimated at the lender-window level and includes lender and event-window fixed effects. Standard errors are clustered by lender.

Equation 2 is the static DiD specification obtained by collapsing the event study into pre- and post-shock periods:

$$Y_{it} = \beta_H \, \text{HighOTD}_i + \beta_P \, \text{Post}_t + \beta_{\text{DiD}} \left(\text{HighOTD}_i \times \text{Post}_t \right) + \mathbf{X}_{it}'\theta + \gamma_s + \delta_t + \varepsilon_{it}, \tag{2}$$

where $Post_t$ equals one for loans originated in July 2021 or later. The vector $\mathbf{X}_{it}$ contains the baseline property controls, including log rentable building area, tenancy type, property taxes per square foot, and local macroeconomic controls. Each property appears once in this specification; t denotes its origination-year cohort. The Post main effect β_P is included separately because loan-origination-year fixed effects cannot absorb it, since the shock falls within calendar year 2021. This DiD specification uses pre- and post-shock periods, and the

binary High-OTD classification is the primary treatment measure throughout the DiD and instrumental-variables analyses.

5.2 Leverage

We begin with leverage, the most direct measure of lending standards. Since our post-shock sample runs just under five years, we split it into eight 15-month periods anchored at the July 2021 shock, four after and three before, and pool everything earlier into one "All Prior" bucket. Figure 2 presents this evidence. Before July 2021, leverage evolves similarly across High-OTD and other lenders. Following the onset of the lending boom, the two series diverge sharply across all three leverage measures, suggesting that the distribution-oriented lending model becomes particularly relevant once lending opportunities expand rapidly.

Figure 3 reports the corresponding event-study estimates. The coefficients before July 2021 are economically small and statistically indistinguishable from zero, providing support for the parallel-trends assumption (formal joint tests in Appendix Table C1). After the onset of the lending boom, the treatment effect becomes positive and remains economically large throughout the post-shock period.

The DiD estimates show that the post-July 2021 leverage gap between High-OTD banks and other lenders increases substantially across all three outcomes, as reported in Table 3. The High-OTD gap in debt per square foot, debt-to-rent, and loan-to-value widens by factors of approximately 4.6 to 5.0 from July 2021 onward.[4]

5.3 Borrower Selection

Lending standards also determine who receives credit. We therefore examine whether High-OTD lenders differentially finance Undiversified Datacenter borrowers. The event study measures this extensive margin at the lender-window level by counting how many Undiver-

[4]The High-OTD gap is the High-OTD-versus-other-lender ratio in each outcome. Because outcomes are in logs, a DiD coefficient of $\hat{\beta}$ implies that this gap changes by a factor of $e^{\hat{\beta}}$ from the pre-boom to the boom period; $e^{1.607} \approx 5.0$, $e^{1.531} \approx 4.6$, and $e^{1.525} \approx 4.6$.

sified Datacenter deals each lender financed. Figure 4 shows this in levels. High-OTD and other lenders follow broadly similar paths before the shock, then diverge afterward; the same figure's Panel B shows the corresponding event study. All earlier windows are pooled into a single coefficient for a stable estimate; the three most recent pre-shock windows and four post-shock windows are shown individually. None of the individual pre-shock coefficients differ significantly from the reference period, consistent with parallel trends (Appendix Table C1); individual post-shock coefficients are noisy given thin per-window cells, but the pooled post-shock estimate below is positive and highly significant.

We then apply an analogous DiD specification, built as a sequence of equal 15-month windows spanning the full sample to ensure balance given some sparsely populated older periods, with window and state fixed effects and state-level macroeconomic controls. As reported in Table 4, the High-OTD gap in Undiversified Datacenter deals financed widens by a factor of approximately 3.3 from July 2021 onward.[5]

We then test whether the leverage premium itself is concentrated among these borrowers. We segment the sample into Undiversified Datacenter and Other borrowers and estimate the same High-OTD $\times$ Post-July 2021 specification (Equation 2) separately for each group. The split-sample estimates reported in Table 5 show that the premium is concentrated among Undiversified Datacenter borrowers. For these borrowers, the High-OTD gap in debt per square foot, debt-to-rent, and loan-to-value widens by factors of approximately 4.4 to 5.3 from July 2021 onward, while none of the corresponding gaps for other borrowers is statistically distinguishable from zero. This split-sample comparison is our strongest evidence against a generic composition story across borrower types. If compositional sorting alone explained the leverage premium, some of it should appear among Other borrowers financed by High-OTD lenders too, yet none does. It does not rule out finer sorting within the Undiversified Datacenter category itself, a concern we address directly with a borrower fixed-effects specification. Together, the borrower-selection and leverage results point to lax lending on

[5]Because the outcome is a Poisson count of deals, a coefficient of $\hat{\beta}$ implies that this gap changes by a factor of $e^{\hat{\beta}}$; $e^{1.179} \approx 3.3$.

both the extensive margin of borrower selection and the intensive margin of loan terms.

These estimates are also robust to extreme observations. Appendix Table C2 shows the estimates remain positive and statistically significant after winsorizing leverage at the 1st and 99th percentiles or trimming observations above the 99th percentile.

6 Robustness and Alternative Explanations

The DiD analysis identifies a leverage premium and a shift in borrower selection associated with distribution-oriented lenders following the datacenter lending boom. This section evaluates whether that finding can be explained by alternative mechanisms rather than weaker lending standards.

No single test resolves identification, so this section addresses five alternative explanations in turn: borrower or project sorting, generic aggressive lending by High-OTD banks, bank size and business-model confounds, post-treatment or endogenous OTD classification, and mechanical construction of leverage measures. We close by comparing our empirical design with that of Purnanandam (2011), the closest study of the same mechanism. Section 7 then addresses the remaining concern that OTD orientation itself may be endogenous, using an instrumental-variables strategy.

6.1 Sorting Across Borrowers or Projects

A further concern is that High-OTD and Low-OTD lenders in the matching sample finance systematically different properties. The matched estimates reported in Appendix Table D1 combine the baseline DiD specification with inverse-probability weighting on common support and coarsened exact matching on state and two property-size bins, the count that maximizes treated-unit coverage under common support. Because matching and inverse-probability weighting require common support between comparably classified lenders, these exercises necessarily restrict the baseline control group to lenders with a defined High- or

Low-OTD classification. This sample restriction is therefore imposed by the matching design rather than chosen as an alternative control group. The resulting High-OTD × Post-Jul'21 estimates remain positive and statistically significant, although somewhat smaller in some specifications, suggesting that observable differences in borrower composition do not explain the baseline findings. This exercise balances properties on observable characteristics, including state and property size, but cannot rule out sorting on unobservable project quality or lender access.

Appendix Table D2 goes a step further on precisely this residual concern. By absorbing borrower fixed effects, it identifies the High-OTD coefficient from each borrower's own pre-to-post change rather than from cross-borrower differences, absorbing persistent borrower characteristics, including unobservable ones; borrowers financed by both a High-OTD and a non-High-OTD lender provide the most direct within-borrower comparison. The leverage premium survives with similar magnitude and remains significant; the wider standard errors that accompany this specification reflect the loss of cross-sectional variation inherent in a same-borrower comparison and are not evidence of a weaker result.

6.2 Generic Aggressive Lending by High-OTD Banks

A further concern is that High-OTD banks simply lend more aggressively in every commercial real estate market. To isolate this possibility, we estimate a pooled cross-sectional specification that abstracts from the timing of the datacenter boom, applied to a separate non-datacenter sample. The placebo estimates reported in Appendix Table D3 show that High-OTD banks extend *less* leverage (Panel A) and do not make larger loans in this market (Panel B). A generic underwriting-aggressiveness explanation therefore cannot account for the main results.

6.3 Bank Size or Business-Model Confounds

Another possibility is that the results simply reflect the behavior of large banks rather than OTD incentives. We therefore control for predetermined bank size using 2021 assets. The leverage and borrower-selection estimates in Appendix Table D4, Panel A, are essentially unchanged. The results are similarly robust to excluding the 15 largest lenders from the sample altogether rather than controlling for size continuously. Leverage and borrower-selection estimates remain consistent (Panel B). Several High-OTD banks well outside the largest lenders in the sample are meaningfully active in this market, including Fifth Third Bank and CIT Group (Appendix Table B3), showing that High-OTD activity is not confined to the largest institutions.

6.4 Post-Treatment or Endogenous OTD Classification

A further concern is that our OTD classification could itself respond to the datacenter lending boom. If a bank's business model shifted because of the boom rather than only causing the leverage response to it, the classification would be contaminated by the very outcome it is meant to predict. We address this by construction. The primary classification uses only bank balance-sheet data from 2017–2020, predetermined relative to the July 2021 onset of the lending boom, so it cannot mechanically reflect any change in bank behavior induced by the boom itself. Panel A of Appendix Table D5 shows the results are robust to rebuilding the classification from a 2022–2025 average instead, applying the identical manual overrides described in Appendix A.3; leverage and borrower-selection estimates remain of similar magnitude and statistically significant, indicating that the specific choice of pre-period is not driving the result, though this alternative-vintage check by construction cannot address concerns about the primary classification's timing itself. Panels B and C of the same table show the results are also stable to excluding Syndicated-type lenders, which are rule-coded as High-OTD rather than scored from Call Report data, and to excluding foreign banks, classified using DealScan syndication behavior rather than Call Report data,

addressing the possibility that these non-standard classification methods are driving the result.

6.5 Mechanical Construction of Leverage Measures and Sample Restrictions

A further concern is that the leverage premium is a mechanical artifact of how leverage is measured, or reflects sample composition, rather than a real underwriting difference. Because loan balances for multi-property loans are allocated across collateral properties by RBA, Appendix Table D6, Panel A, restricts the sample to single-property loans, eliminating any role for the allocation rule entirely; Panel B instead allocates loan balances equally rather than proportionally across collateral properties. Panel C drops properties classified by CoStar as abandoned, proposed, final planning, or under construction, which may not yet be, or may no longer be, operational. The estimates in all three panels remain positive and statistically significant for both leverage and borrower selection.

6.6 Comparison with Purnanandam (2011)

Because Purnanandam (2011) provides the closest empirical design for studying OTD incentives, we compare our identification strategy directly with his. Both designs combine a market-wide change in lending opportunities with pre-existing differences in lenders' distribution orientation, but differ in the source of that change, the measurement of distribution orientation, and the outcomes observed.

Purnanandam (2011) studies the 2007 collapse of the secondary mortgage market, an abrupt contraction in banks' ability to distribute residential mortgages. We instead study the AI-driven expansion in datacenter lending, with July 2021 marking the empirical onset of the lending boom in our data. His measure of OTD behavior comes directly from HMDA loan-sale data, whereas we measure distribution orientation using predetermined

bank balance-sheet characteristics because no comparable loan-level disclosure exists for commercial datacenter mortgages. Finally, Purnanandam observes ex-post loan performance, whereas we study underwriting decisions at origination through leverage and borrower selection. Appendix Table D7 summarizes these differences. The two designs therefore provide complementary evidence on the same incentive mechanism across different changes in lending opportunities and at different stages of the credit cycle.

7 Instrumental-Variables Analysis

Our baseline identification strategy exploits the differential response of High-OTD and other lenders to the datacenter lending boom. Although the event-study evidence supports the baseline DiD specification, and Section 6 shows the results are robust to a battery of alternative explanations, OTD orientation may itself be endogenous to underwriting behavior. We therefore complement the DiD design with an instrumental-variables strategy that exploits predetermined cross-bank variation in lending intensity measured before the boom, addressing this remaining component of the identification challenge.

7.1 Identification

We instrument High-OTD status using each bank's predetermined 2020 loan-to-assets ratio. The intuition is that banks with smaller loan portfolios relative to assets before the datacenter lending boom were already more distribution-oriented and thus more likely to be classified as High-OTD, while the loan-to-assets ratio itself was determined before either the boom or the subsequent monetary tightening cycle. The strongest evidence for this exclusion restriction is that the instrument has no relationship with leverage in a placebo window predating the boom, discussed below.

The instrumental-variables analysis is estimated on the property-level sample, where properties are indexed by i and lending banks by b. The first-stage specification is

$$\text{HighOTD}_b = \pi Z_b + \mathbf{X}'_{i(b)t}\theta + \gamma_s + \delta_t + \nu_{i(b)t}, \tag{3}$$

where Z_b denotes bank b's 2020 loan-to-assets ratio and $\mathbf{X}_{i(b)t}$, γ_s, and δ_t are defined as in Equation 2. The corresponding second-stage specification is

$$Y_{it} = \beta \widehat{\text{HighOTD}}_{b(i)} + \mathbf{X}'_{it}\theta + \gamma_s + \delta_t + \varepsilon_{it}, \tag{4}$$

Equations 3–4 instrument the cross-sectional level of High-OTD status.

The exclusion restriction is that the 2020 loan-to-assets ratio affects post-shock underwriting outcomes only through its effect on High-OTD status and is otherwise unrelated to pre-existing lender specialization, funding costs, or anticipation of the datacenter boom. Because the instrument is measured in 2020, before the 2021 acceleration in datacenter lending and before the subsequent monetary tightening cycle, it is predetermined with respect to both shocks.[6]

Because the exclusion restriction cannot be tested directly, we evaluate its plausibility through a series of validation exercises. Figure 5 shows the headline evidence. Panel A confirms a strong first stage, with the instrument sharply predicting High-OTD status. Panels B and C show the corresponding timing break in the reduced-form relationship between the instrument and ln(DebtPSF). The relationship holds in the real post-boom window but is statistically indistinguishable from zero in a placebo window that predates the datacenter lending boom. Because the instrument is predetermined, this timing break is inconsistent with a relationship that was already operating in the pre-boom period or with the instrument simply capturing an immediate response to the datacenter lending boom itself. Appendix Figure E1 provides further support, showing that lenders classified by the instrument evolve similarly before the datacenter lending boom and diverge only afterward.

The Appendix provides additional supporting evidence. Appendix Table E2 shows the

[6]The loan-to-assets ratio is highly persistent across years, with cross-bank correlations of 0.873 to 0.966 over 2021–2025 (Appendix Table E1).

instrument is not significantly related to pre-boom market share; pre-boom origination volume is significantly lower for banks with a higher, Low-OTD-predicting loan-to-assets ratio, meaning banks predicted to be High-OTD were, if anything, already somewhat more active in datacenter lending before the boom. This volume result is therefore not on its own favorable evidence against pre-existing specialization; the null market-share result, and more directly the null pre-boom leverage placebo below, are the stronger defenses of the exclusion restriction. Panel C of the same table tests for pre-boom differences in bank quality. Return on assets (ROA), unlike net interest margin, nets out funding costs and is not mechanically tied to a bank's loan-versus-trading-asset mix, making it the more informative test of whether the instrument proxies for broader bank quality or risk appetite rather than loan-heavy composition alone. Predicted-Low-OTD banks show no significant difference in pre-boom ROA, providing no evidence for that alternative channel; the significant NIM difference, by contrast, is consistent with the instrument simply reflecting loan-share composition. Appendix Table E3 demonstrates that the results are robust to controlling for alternative funding structures. Appendix Table E4, Panel A, formally tabulates the placebo-window test underlying the timing break above, and Panel B reports alternative-instrument-vintage tests; neither shows evidence of spurious pre-existing trends. Although these exercises cannot establish the exclusion restriction, they provide supporting evidence for the credibility of the identifying assumptions underlying the IV design.

7.2 Results

IV estimates of the relationship between High-OTD status and borrower leverage are reported in Table 6. The cross-sectional IV (Equation 4) instruments High-OTD status with a strong first stage (Kleibergen–Paap $F = 25.05$), supporting the interpretation that the leverage premium is not driven solely by endogenous selection into High-OTD status.[7]

[7]As an untabulated robustness check, not the specification reported in Table 6, we also estimate a DiD-IV specification that additionally instruments $\text{HighOTD}_b \times \text{Post}_t$ with $Z_b \times \text{Post}_t$, adding $\beta_P \,\text{Post}_t + \beta_{\text{DiD}}\big(\widehat{\text{HighOTD}_b} \times \text{Post}_t\big)$ to Equation 4. This yields closely aligned estimates, with the expected weaker

The cross-sectional IV extends to the extensive-margin borrower-selection outcome as well. The cross-sectional 2SLS estimate remains positive and significant with a strong first stage (KP-$F = 35.65$), supporting the interpretation that this margin is not driven solely by endogenous selection into High-OTD status either (Appendix Table E5).[8]

8 Mechanism: Loan Size

Why do High-OTD lenders extend substantially more leverage? We argue that the answer lies in loan size. The evidence below shows that the post-boom increase in High-OTD lending is not accompanied by a corresponding increase in the size of financed properties; instead, High-OTD lenders extend much larger loans. This distinction is important because, in a market where the number of financeable projects is limited, increasing loan size is the primary way to expand origination volume, and hence the fee revenue generated from originating and distributing credit.

Figure 6 provides the descriptive evidence. Loan balances expand much more rapidly than the physical scale of financed properties. Figure 7, Panel A, shows that loan amounts evolve similarly across High-OTD and other lenders before July 2021 but diverge sharply afterward. Panel B confirms this pattern in the corresponding event-study specification. Table 7 formalizes this comparison. The High-OTD gap in loan size widens by approximately 4.8-fold from July 2021 onward, while there is no statistically detectable increase in financed property RBA.[9] High-OTD lenders do finance somewhat larger properties on average, about

first stage for an interacted instrument (KP-$F = 11.06$ for the debt-per-square-foot specification).

[8]This result is also stable across analogous funding-control, pre-boom placebo, and instrument-vintage checks (Appendix Table E6), which mirror the leverage versions of these checks reported in Appendix Tables E3 and E4. As a further untabulated robustness check, a DiD-IV specification applies the same extension as the leverage DiD-IV above, instrumenting $\text{HighOTD}_b \times \text{Post}_t$ with $Z_b \times \text{Post}_t$. This borrower-selection estimate is similarly positive and significant, with a weaker first stage (KP-$F = 17.81$), the same tradeoff as an interacted instrument.

[9]The loan-size premium is stable in sign, magnitude, and significance across the full robustness battery used for the leverage and borrower-selection results, including winsorizing and trimming, excluding non-operational properties, single-property loans only, equal-split allocation, a predetermined bank-assets control, propensity-score and exact matching, excluding syndicated loans, excluding foreign banks, and rebuilding the High-OTD classification from a 2022–2025 average instead of the primary 2017–2020 window.

26% larger, a premium specific to the datacenter market that does not appear in the non-datacenter placebo sample (Appendix Table D3, Panel B); as shown above, this size gap does not widen following July 2021 (Table 7), even as the loan-size gap widens sharply. The increase in leverage therefore reflects larger loans rather than larger underlying assets.[10] Concentration is starkest among the very largest transactions.[11]

Datacenter project finance differs from residential mortgage lending in that lending opportunities are scarce. Only 438 properties receive mortgage financing after the onset of the lending boom, and each transaction involves substantial fixed costs, including legal, financial, and technical due diligence on specialized power, cooling, and network infrastructure, negotiation of hyperscale lease contracts, and underwriting collateral in a market with no prior default cycle. We interpret the loan-size premium documented above as evidence consistent with volume-maximizing behavior under OTD incentives, rather than as evidence that High-OTD lenders simply finance larger properties.

The same acceleration in loan size appears independently in DealScan syndicated-loan data, outside our mortgage sample entirely (Appendix Figure F1). The non-datacenter placebo reinforces this interpretation. High-OTD banks exhibit neither a loan-size premium nor unusually large financed properties outside the datacenter market (Appendix Table D3, Panel B), where the supply of lending opportunities is not similarly constrained.

Finally, the loan-size premium is consistent with how distribution-oriented lenders are compensated. Syndicated loan fees, including arranger, commitment, and participation fees, are typically assessed as a percentage of facility amount rather than as a fixed payment per transaction (Berg et al., 2016). Two disclosed datacenter credit agreements illustrate this structure. Each recurring fee is stated as a percentage of committed or outstanding exposure (Appendix Table F1). Because individual datacenter facilities can reach billions of dollars, increasing loan size substantially increases fee income, consistent with fee structures tied to

[10]Loan size is observed directly from origination records and allocated across collateral properties in proportion to RBA for multi-property loans, as described in Appendix A.5.

[11]Of the 42 loans in the sample at $1 billion or more, 41 (98%) are High-OTD-financed; of the 83 loans at $500 million or more, 73 (88%) are.

loan volume in an OTD business model.

The resulting leverage has important implications for credit risk. Measured against real-estate collateral alone, required revenue at the most leveraged facilities exceeds observed market rents (Appendix Figure F2), and over a third of properties carry debt exceeding their estimated real-estate collateral value (Appendix Table A1). These patterns leave highly leveraged borrowers with limited collateral cushions against a decline in datacenter revenues and depreciating IT assets, making continued demand for datacenter capacity central to debt repayment.

9 Conclusions

This paper brings into view a previously opaque segment of the AI financing ecosystem, property-level datacenter mortgages. We build a new property-level database linking U.S. datacenter mortgages to lender balance sheets and use it to document a large, rapidly expanding, and highly concentrated market. The expansion is dominated by distribution-oriented lenders and directed primarily toward Undiversified Datacenter borrowers, making this market a natural setting in which to study whether OTD incentives shape lending standards.

The abrupt onset of the datacenter lending boom in July 2021 creates a sharp break for a difference-in-differences design comparing properties financed by High-OTD lenders with those financed by other lenders. Across multiple identification checks, High-OTD lenders increase leverage substantially and disproportionately finance Undiversified Datacenter borrowers, with larger loan sizes driving the leverage premium. Together, these patterns are consistent with distribution incentives shaping both the scale and allocation of credit.

The broader concern is that OTD incentives are emerging in a market whose risk-management infrastructure has yet to be tested under stress. Datacenter lending has developed rapidly, without the benefit of a prior default cycle to establish recovery values,

underwriting benchmarks, or investor discipline. Our findings have direct implications for credit risk: distribution-oriented lenders relax underwriting standards, extending more leverage and channeling credit toward riskier, less diversified borrowers, through larger loan sizes, in a new and rapidly growing asset class that has not yet faced a downturn. This provides a basis for supervisory attention while the market is still forming, rather than after distress has already been realized.

Several questions remain open for future research. Where does the credit risk associated with these mortgages ultimately reside after origination and distribution? How will these loans perform through the market's first default cycle, and will High-OTD lenders differ from other lenders in workout behavior? How will technological obsolescence affect collateral values as specialized IT equipment ages? Future work could also examine how equity structures, colocation contracts, vendor guarantees, and lease terms interact with credit risk; whether similar lending patterns emerge outside the United States; and how a shift toward hyperscaler-owned rather than leased capacity would affect financing volumes and borrower risk.

References

Acharya, V. V., Schnabl, P., and Suarez, G. (2013). Securitization without risk transfer. *Journal of Financial Economics*, 107(3):515–536.

Adrian, T. and Shin, H. S. (2010). Liquidity and leverage. *Journal of Financial Intermediation*, 19(3):418–437.

Akerlof, G. A. (1970). The market for "lemons": Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, 84(3):488–500.

Aobdia, D., Dou, Y., and Kim, J. (2021). Public audit oversight and the originate-to-distribute model. *Journal of Accounting and Economics*, 72(1):101420.

Berg, T., Saunders, A., and Steffen, S. (2016). The total cost of corporate borrowing in the loan market: Don't ignore the fees. *Journal of Finance*, 71(3):1357–1392.

Berndt, A. and Gupta, A. (2009). Moral hazard and adverse selection in the originate-to-distribute model of bank credit. *Journal of Monetary Economics*, 56(5):725–743.

Black, L., Chu, C. S., Cohen, A., and Nichols, J. B. (2012). Differences across originators in CMBS loan underwriting. *Journal of Financial Services Research*, 42(1):115–134.

Bord, V. M. and Santos, J. A. C. (2015). Does securitization of corporate loans lead to riskier lending? *Journal of Money, Credit and Banking*, 47(2–3):415–444.

CoreSite, L.P. (2019). Amended and restated credit agreement. U.S. Securities and Exchange Commission, Exhibit 10.1.

Darmouni, O., Lehner, C., and Zhang, Y. (2025). The energy cost of AI and data centers. Working Paper, SSRN 6751399.

Dell'Ariccia, G., Igan, D., and Laeven, L. (2012). Credit booms and lending standards: Evidence from the subprime mortgage market. *Journal of Money, Credit and Banking*, 44(2–3):367–384.

DeMarzo, P. and Duffie, D. (1999). A liquidity-based model of security design. *Econometrica*, 67(1):65–99.

Gargano, A. and Giacoletti, M. (2025). Subsidizing the Cloud: U.S. State Incentives to Data Centers. Working paper.

Gorton, G. and Pennacchi, G. (1995). Banks and loan sales: Marketing nonmarketable assets. *Journal of Monetary Economics*, 35(3):389–411.

Gustafson, M. T., Ivanov, I. T., and Meisenzahl, R. R. (2021). Bank monitoring: Evidence from syndicated loans. *Journal of Financial Economics*, 139(2):452–477.

Heiss, A., Sautner, Z., and Schmid, T. (2026). Powering AI: How Do Data Centers Affect Renewable Energy Investment? Swiss Finance Institute Research Paper No. 26-29.

Ivashina, V. and Sun, Z. (2011). Institutional demand pressure and the cost of corporate loans. *Journal of Financial Economics*, 99(3):500–522.

Keys, B. J., Mukherjee, T., Seru, A., and Vig, V. (2010). Did Securitization Lead to Lax Screening? Evidence from Subprime Loans. *The Quarterly Journal of Economics*, 125(1):307–362.

Loutskina, E. (2011). The role of securitization in bank liquidity and funding management. *Journal of Financial Economics*, 100(3):663–684.

Loutskina, E. and Strahan, P. E. (2009). Securitization and the declining impact of bank finance on loan supply: Evidence from mortgage originations. *The Journal of Finance*, 64(2):861–889.

McGowan, D. and Nguyen, H. (2023). To securitize or to price credit risk? *Journal of Financial and Quantitative Analysis*, 58(1):289–323.

McGowan, D., Nguyen, H., and Schaeck, K. (2024). Deposit competition and mortgage securitization. *Journal of Money, Credit and Banking*.

Mian, A. and Sufi, A. (2009). The Consequences of Mortgage Credit Expansion: Evidence from the U.S. Mortgage Default Crisis. *The Quarterly Journal of Economics*, 124(4):1449–1496.

Phoenix Data Center Acquisitions LLC (2024). Revolving credit agreement. U.S. Securities and Exchange Commission, Exhibit 10.11.

Piskorski, T., Seru, A., and Vig, V. (2010). Securitization and distressed loan renegotiation: Evidence from the subprime mortgage crisis. *Journal of Financial Economics*, 97(3):369–397.

Puri, M., Rajan, R., and Singh, I. (2011). On the benefits of concurrent lending and underwriting. *The Journal of Finance*, 66(4):1187–1243.

Purnanandam, A. (2011). Originate-to-distribute model and the subprime mortgage crisis. *The Review of Financial Studies*, 24(6):1881–1915.

Stiglitz, J. E. and Weiss, A. (1981). Credit rationing in markets with imperfect information. *The American Economic Review*, 71(3):393–410.

Sufi, A. (2007). Information asymmetry and financing arrangements: Evidence from syndicated loans. *The Journal of Finance*, 62(2):629–668.

Van Nieuwerburgh, S. (2026). Financing the AI buildout. Columbia Business School Working Paper.

Yue, D. and Zeng, Y. (2026). The local economic effects of data center entry. Working paper.

Figures

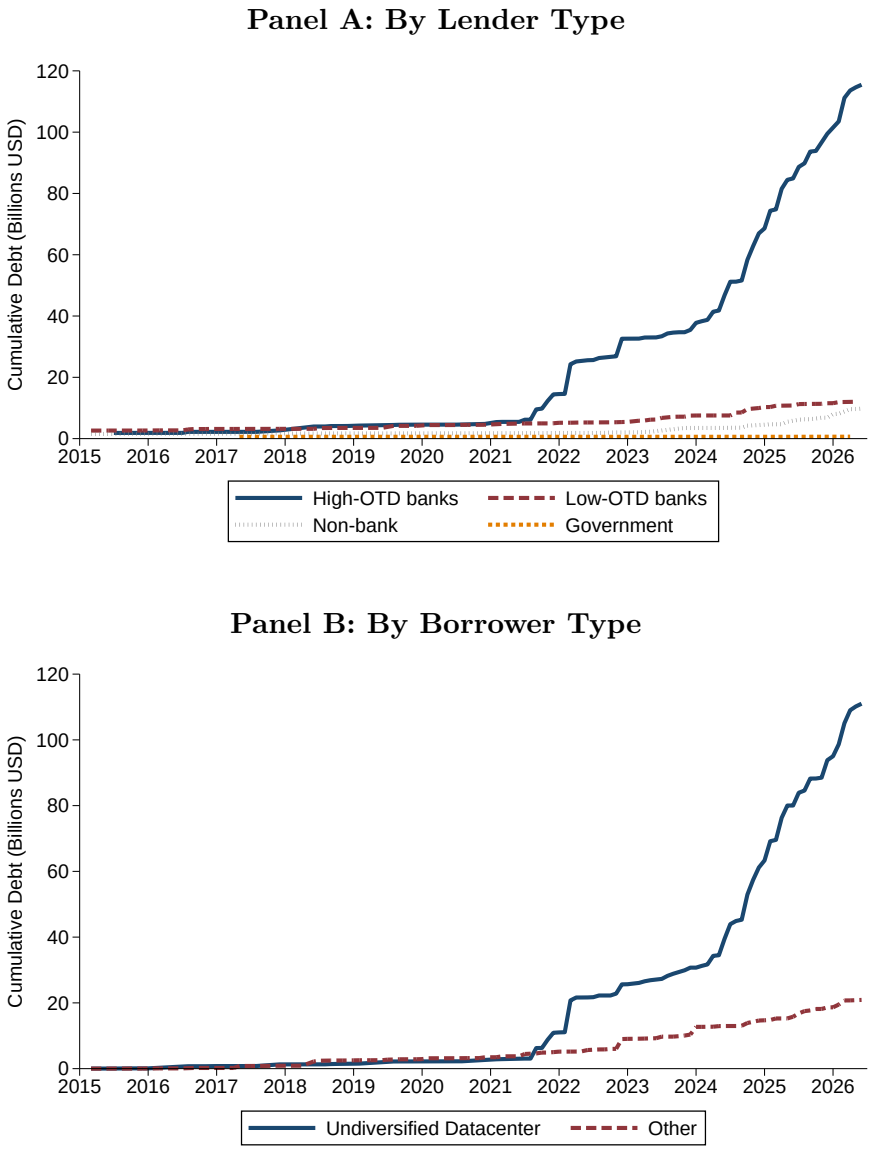


Figure 1. Cumulative Datacenter Mortgage Origination

Loan amounts summed cumulatively by origination year. OTD denotes originate-to-distribute. Panel B distinguishes Undiversified Datacenter borrowers, whose revenues depend primarily on datacenter leasing and who serve as our proxy for the borrowers most exposed to AI-demand risk.

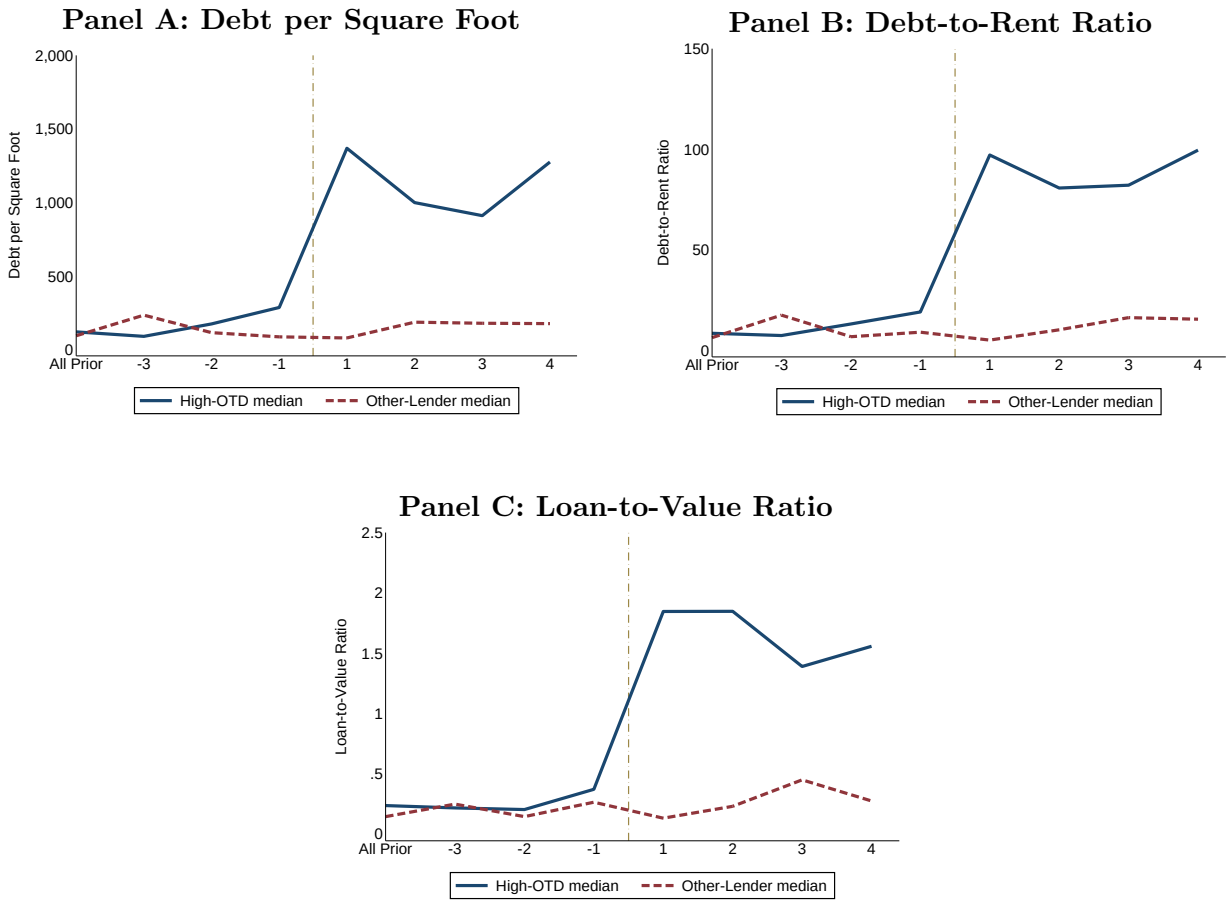


Figure 2. Leverage Levels Before and After the Shock

Median leverage of transactions by lender type. The shock is anchored at July 2021; each period is 15 months, before and after the shock, except "All Prior," which pools every origination before the first period.

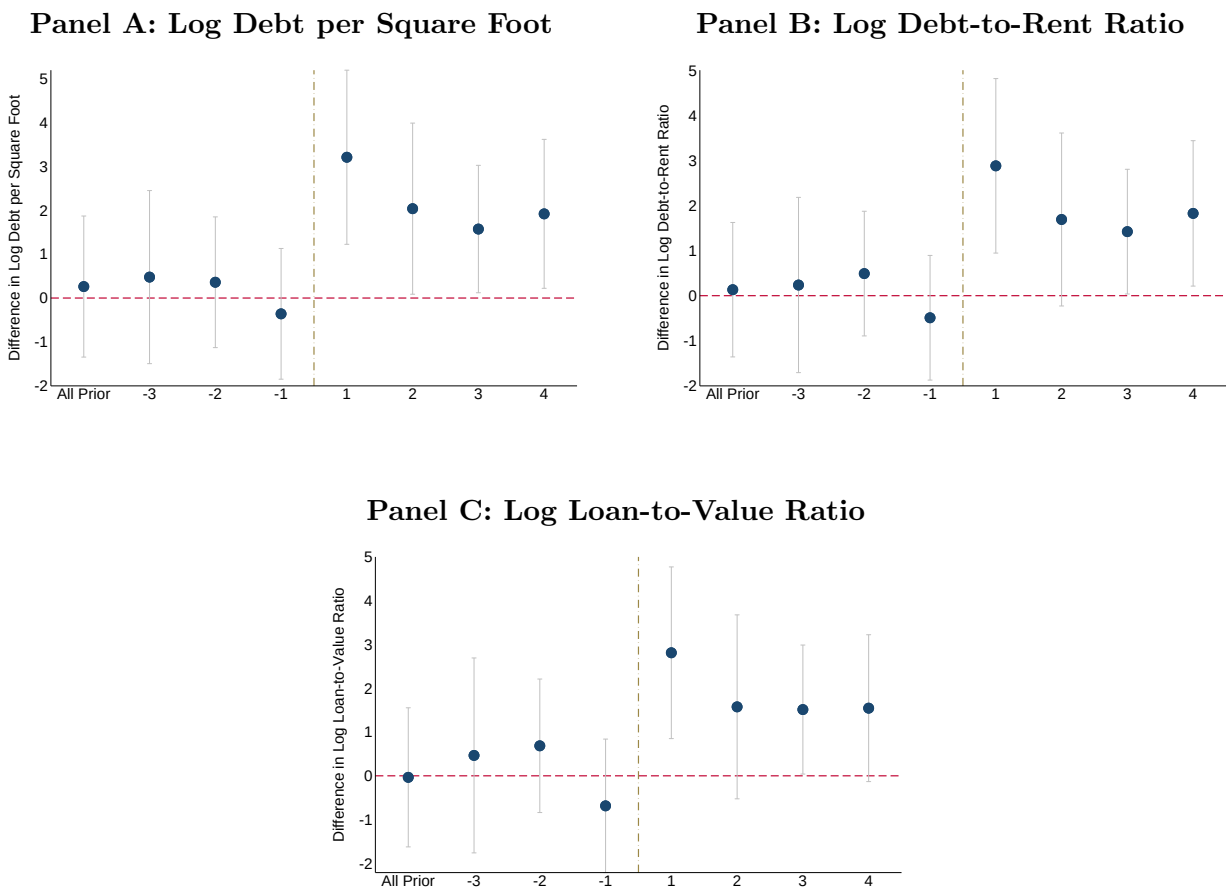
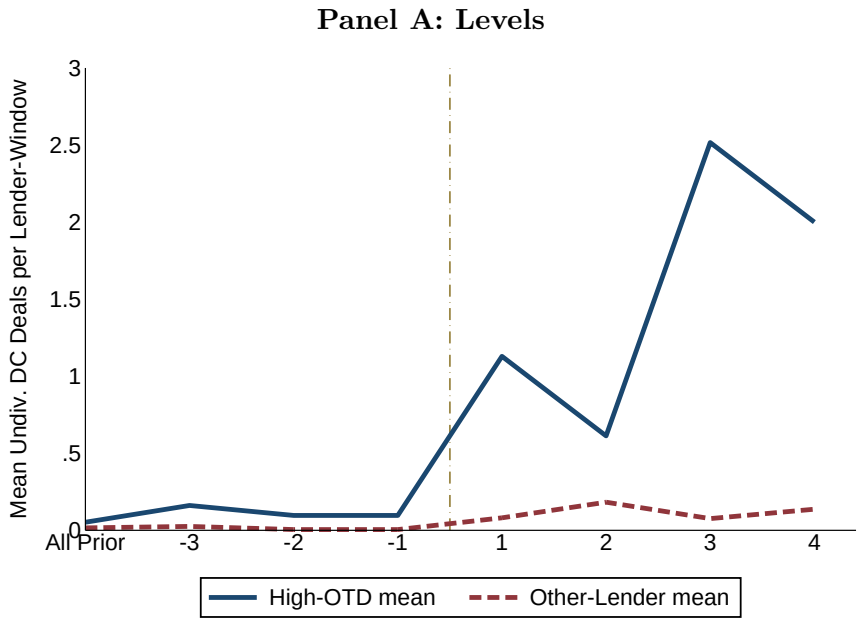


Figure 3. Leverage Rises After the Boom

Event-study estimates compare lending standards at High-OTD and other lenders, in the same periods as Figure 2, following Equation 1. The specifications include state, origination-year, and event-window fixed effects. Error bars are 95% confidence intervals.



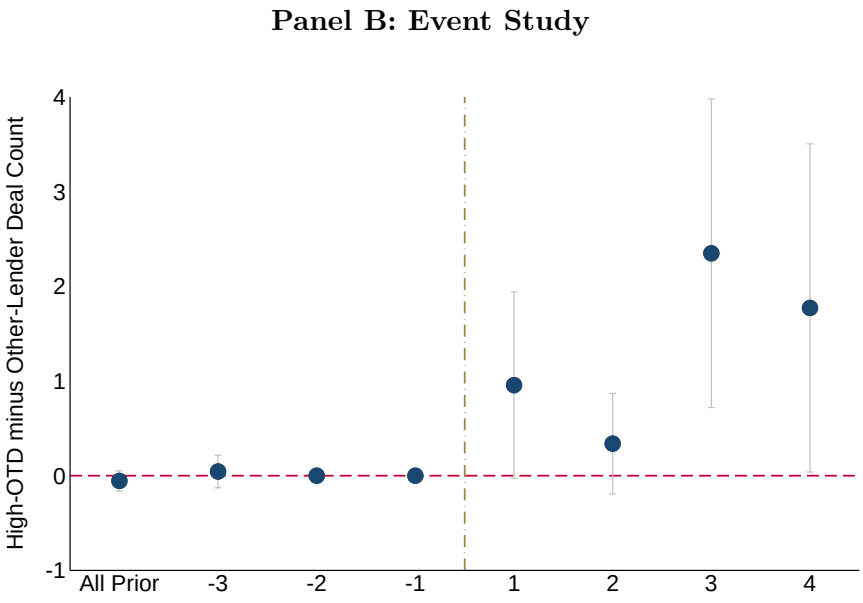


Figure 4. Financing of Undiversified Datacenter Borrowers

Panel A reports mean Undiversified Datacenter deals financed per lender-window, by lender type, in the same periods as Figure 2. Panel B reports the corresponding event-study estimates, with window and lender fixed effects. Error bars are 95% confidence intervals.

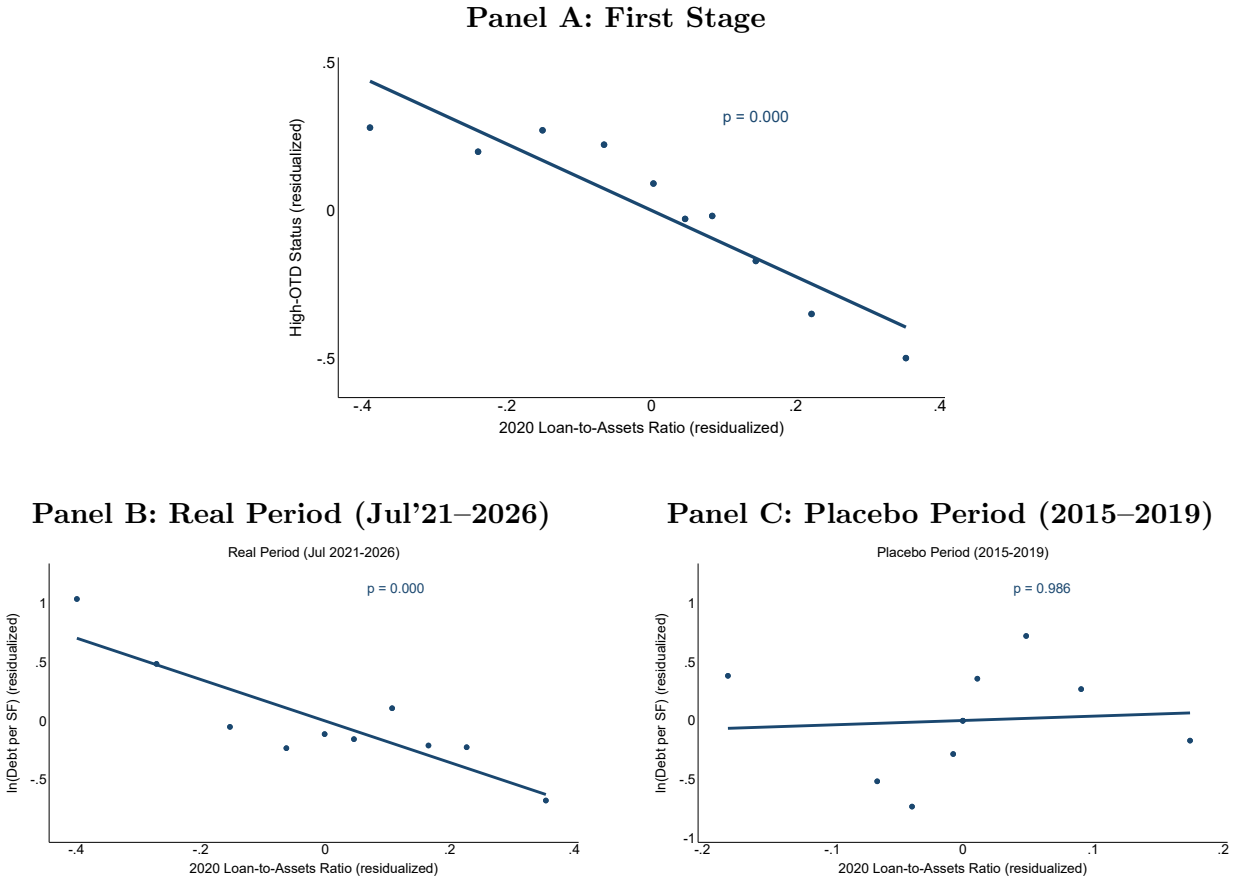


Figure 5. Instrument Validity: Relevance and Placebo Timing

Panel A plots the binary High-OTD indicator against the bank's predetermined 2020 loan-to-assets ratio, both residualized on the controls in Equation 2 (10 deciles). Panels B and C plot the reduced-form relationship between the instrument and ln(DebtPSF), residualized on the same controls, estimated separately in the real post-boom window (Panel B) and in a placebo window that predates the datacenter lending boom (Panel C); both panels share the same y-axis scale.

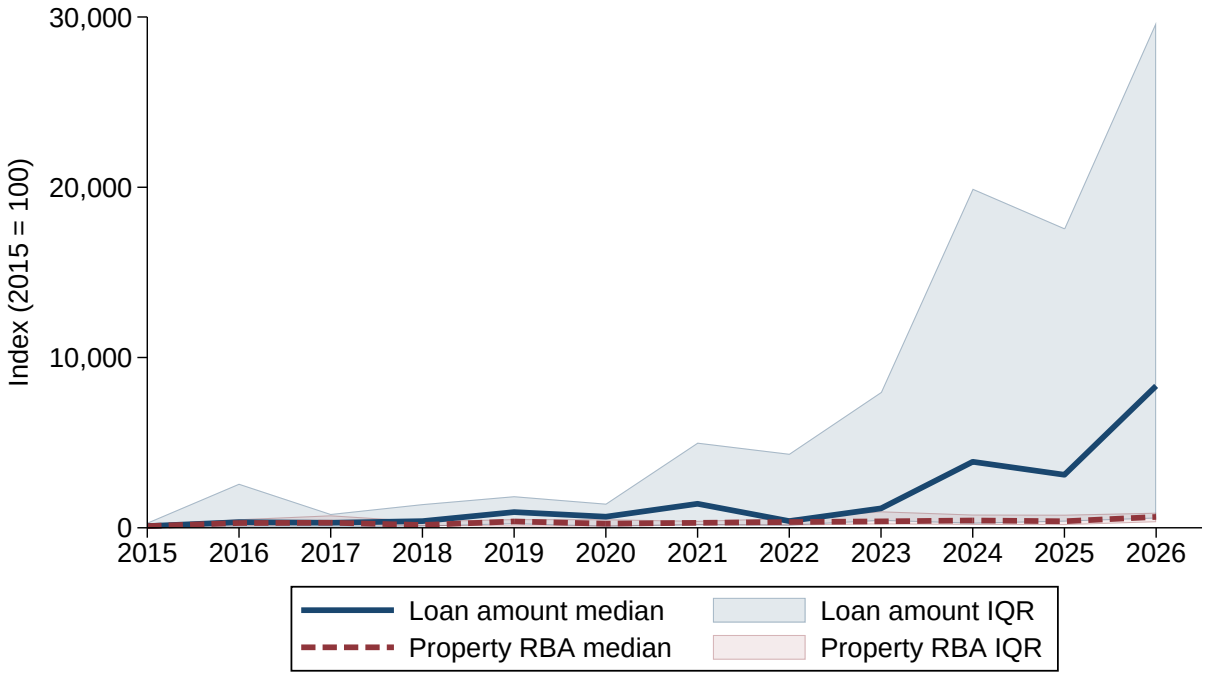


Figure 6. Loan and Property Size

Median loan amount and median rentable building area (RBA), each indexed to its 2015 level (=100).

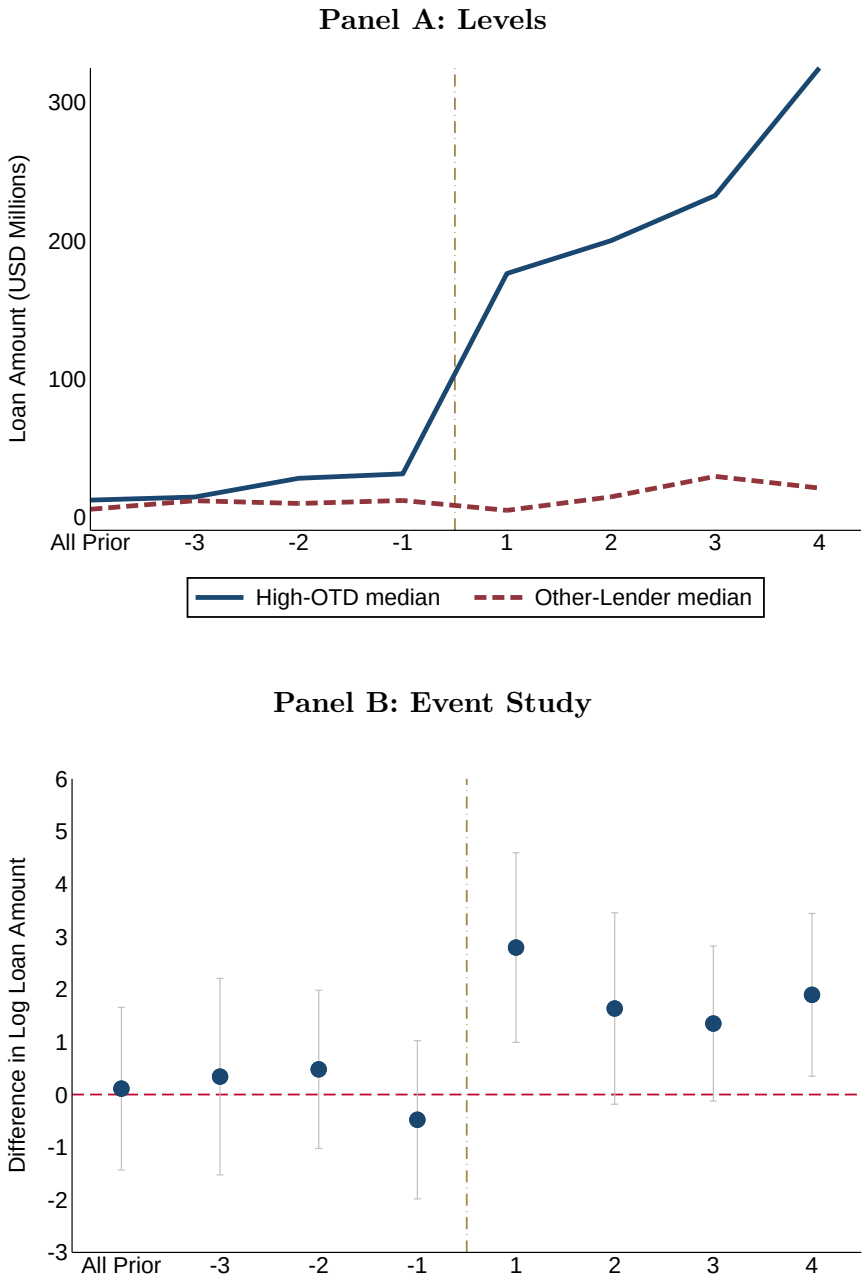


Figure 7. The High-OTD Loan-Size Gap Rises After the Boom

Panel A reports median loan amount ($M) by lender type, in levels, in the same periods as Figure 2. Panel B reports the corresponding event-study estimates of log loan amount, following Equation 1, with state, origination-year, and event-window fixed effects. Error bars are 95% confidence intervals.

Tables

Table 1. Summary Statistics

Variable	N	Median	SD	Min	Max
Panel A: Property characteristics					
Rentable building area (SF)	668	140,000	271,714	7,920	3,000,000
Rent per square foot (USD)	668	12.8	9.0	2.3	122.9
Taxes per square foot (USD)	668	2.2	3.8	0.0	29.5
Years since origination	668	4.0	6.9	0.0	34.0
Panel B: Leverage outcomes					
Loan amount (USD M)	668	33.5	389.7	0.0	3,500.0
Debt per square foot (USD)	668	311	3,464	0.0	31,867
Debt-to-rent ratio	668	22.9	325.7	0.0	3,561
Loan-to-value ratio	668	0.4	7.4	0.0	73.9

Note: Years since origination is measured as of 2026. For multi-property loans, loan amounts are allocated across collateral properties in proportion to RBA. See Appendix A.5 for variable construction.

Table 2. Borrowers, Lenders, and the OTD Classification

Panel A: Borrower type

	Undiversified Datacenter	Other	Difference
Properties	397	271	126
Median loan amount (USD M)	64.0	22.1	41.9
Median property RBA (SF)	175,000.0	109,450.0	65,550.0
Median debt per square foot (USD)	440.7	187.9	252.8

Panel B: Lender type

	High-OTD	Low-OTD	Difference
Properties financed	332	202	130
Median loan amount (USD M)	148.7	10.0	138.7
Median property RBA (SF)	161,901.5	101,000.0	60,901.5
Median debt per square foot (USD)	844.2	128.0	716.2

Panel C: Bank balance sheet characteristics, 2017–2020 (%)

	High-OTD	Low-OTD	Difference
Loan assets/assets	50.97	71.72	−20.75***
Repo/assets	10.66	0.42	10.24**
Trading assets/assets	4.44	0.11	4.33***

Note: Panels A and B report property-level counts and medians. Panel B excludes non-bank and government lenders. Panel C reports bank-level group means of the regulatory ratios underlying the OTD score in Equation A1, restricted to banks with a matched Call Report score. Differences in Panel C use unequal-variance tests of equality. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 3. The OTD Leverage Premium

	ln(DebtPSF)	ln(DebtToRent)	ln(LTV)
High-OTD × Post-Jul'21	1.607***	1.531***	1.525***
	(0.409)	(0.413)	(0.438)
Controls and FE	Yes	Yes	Yes
R^2	0.401	0.382	0.363
Observations	668	668	668

Note: All specifications estimate Equation 2 and include controls for property size, local macro conditions, tenancy type, and taxes per square foot, as well as state and loan-origination year fixed effects. Standard errors clustered at the lender level are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 4. OTD and Borrower Selection

	Undiv. DC Count
High-OTD × Post-Jul'21	1.179***
	(0.425)
Controls and FE	Yes
Observations	10,556

Note: This table reports estimates from a DiD specification analogous to Equation 2, estimated as a Poisson count model at the lender-window-state level. The outcome is the number of Undiversified Datacenter deals financed by each lender in each state and 15-month window. The panel spans all 668 properties in the main sample, collapsed to lender-window-state cells with zero-fill within each lender's revealed footprint. The specification includes window and state fixed effects and state-level macroeconomic controls. Standard errors are clustered by lender. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 5. The Leverage Premium Concentrates in Undiversified Datacenter

	ln(DebtPSF)	ln(DebtToRent)	ln(LTV)
Panel A: Undiversified Datacenter borrowers			
High-OTD × Post-Jul'21	1.673***	1.480**	1.584**
	(0.607)	(0.632)	(0.640)
R^2	0.438	0.431	0.394
Observations	397	397	397
Panel B: Other borrowers			
High-OTD × Post-Jul'21	1.056	1.007	0.911
	(0.702)	(0.731)	(0.788)
R^2	0.483	0.439	0.437
Observations	271	271	271
Controls and FE	Yes	Yes	Yes

Note: This table reports split-sample estimates of the baseline DiD specification (Equation 2), splitting the sample into Undiversified Datacenter and Other borrowers and estimating each panel separately. Controls and fixed effects are as in Table 3. Standard errors clustered at the lender level are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 6. IV Estimates

	ln(DebtPSF)	ln(DebtToRent)	ln(LTV)
OLS	0.739**	0.844***	0.766***
	(0.287)	(0.290)	(0.282)
2SLS: Cross-sectional IV	1.309***	1.427***	1.234***
	(0.401)	(0.394)	(0.366)
Controls and FE	Yes	Yes	Yes
Observations	475	475	475
KP-F	25.05	25.05	25.05

Note: This table reports estimates from the cross-sectional IV specification (Equations 3–4), restricted to bank lenders; OLS uses the bank-only subsample for direct comparability. The excluded instrument is measured in 2020, before the sharp acceleration in datacenter lending beginning in 2021. Standard errors are clustered at the lender level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 7. Loan Scale and Distribution Incentives

	ln(Loan Amount)	ln(RBA)
High-OTD × Post-Jul'21	1.573***	−0.075
	(0.424)	(0.196)
Controls and FE	Yes	Yes
R^2	0.411	0.297
Observations	668	668

Note: This table reports estimates from the baseline DiD specification (Equation 2). Loan amount and RBA are outcomes rather than controls, since deal scale is the outcome being tested. Standard errors clustered at the lender level are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Appendix A: Data and Variable Construction

This appendix describes the data sources, cleaning procedures, and construction of key variables used in the analysis.

A.1 Data Sources and Sample Construction

CoStar datacenter inventory. The starting point is the CoStar national inventory of U.S. datacenter properties as of May 2026. The unit of observation in the main analysis is a property, identified by CoStar. For each property, CoStar reports physical characteristics (rentable building area, year built, and development status), location (state, county, market and submarket), leasing outcomes (shell rent and percent leased), and tax-related fields. Table 1 reports summary statistics for the resulting property universe.

Loan origination records. CoStar also provides loan origination information for a subset of properties, including origination amount, origination date, loan type, interest rate fields, and the originating lender name. The loan information is used to construct property-level leverage outcomes. Because a single loan may be collateralized by multiple properties, loan balances are allocated across collateral properties prior to constructing per-square-foot measures (see Appendix A.5).

Lender OTD measures. We link lenders to regulatory Call Report data to obtain OTD-based distribution orientation measures; construction of the OTD score is described in Appendix A.3.

DealScan syndicated loan data. DealScan is used in two distinct ways. First, we match DealScan lender-tranche records to our own lenders and examine their syndication behavior across all of their DealScan facilities, not only datacenter-related ones, to validate the OTD classification (Appendix A.3). Second, a separate borrower-level sample of datacenter-related facilities confirms loan-size trends outside our mortgage data (Appendix Figure F1). In this sample, Undiversified Datacenter borrowers include firms whose primary business is

datacenter ownership and operation, the sample includes all syndicated facilities, not only mortgage or real estate debt, and the comparison group comprises non-datacenter industrial tranches coded as real estate or construction loans.

Local macro controls. County-level annual GDP, population, and unemployment rate come from the Bureau of Economic Analysis and Bureau of Labor Statistics, compiled in a single macro data file and merged to each property by county and origination year. All three enter regressions in logs (ln(GDP), ln(Population), ln(Unemployment Rate)). A small number of Virginia counties with missing GDP, an artifact of BEA's treatment of independent cities, are imputed using the nearest county with non-missing data.

Data vintages. The CoStar datacenter inventory was accessed in May 2026 and reflects the property universe as of that date. FFIEC Call Report data were obtained via the Federal Reserve Bank of New York in January 2026 and cover 1959 through 2025 Q4. DealScan data were obtained in July 2026 and cover facilities through March 2026. All data vintages are fixed for replication purposes.

A.2 Estimation Samples

Our baseline empirical analysis uses the full sample of 668 mortgaged datacenter properties. Some specifications necessarily rely on different samples because of data availability or the empirical design. The instrumental-variables analysis is restricted to bank lenders that can be matched to Call Report data. Borrower heterogeneity analyses split the sample by borrower type. The borrower-selection analysis aggregates the data to lender-by-state-by-time-window observations, while some robustness exercises impose additional restrictions, such as limiting the sample to single-property loans or excluding non-operational properties. The non-datacenter placebo (Appendix Table D3) uses a separate CoStar extract of non-datacenter industrial commercial real estate loans, entirely outside the datacenter mortgage sample. Unless otherwise stated, results are estimated on the largest sample available for the corresponding empirical exercise.

A.3 Borrower and Lender Classification

Borrower types. We classify borrowers into four groups: Undiversified Datacenter, Big Tech, Government, and Other. The classification uses standardized borrower names and borrower group labels derived from CoStar ownership fields. The indicator `UndiversDC` equals one for Undiversified Datacenter properties.

Extensive margin (deal count). For Table 4, we collapse the sample to lender-by-window-by-state cells, where windows are equal 15-month periods spanning the full sample. Within each cell, the outcome is the count of Undiversified Datacenter deals financed, $\text{ExtCount}_c = \sum_{i \in c} \text{UndiversDC}_i$, where c indexes a lender $\times$ window $\times$ state cell, zero-filled within each lender's revealed state footprint. For the corresponding event study (Figure 4), we aggregate counts across states to the lender-window level and include lender and window fixed effects. This measures how many such deals each lender finances while avoiding comparisons driven by differences in lenders' geographic footprints. Earlier windows are pooled into a single coefficient, and the three most recent pre-shock windows and four post-shock windows are estimated individually.

Lender mapping and lender types. The loan records report a lender name (originator). We standardize lender names and match them via a name crosswalk to Call Report regulatory identifiers; High/Low-OTD status then follows mechanically from the decile rule described below. Non-bank and government/cooperative lenders are assigned separately by lender type. See Foreign banks, discussed below, for lenders that cannot be matched to Call Report data or whose match is unrepresentative.

Originate-to-distribute (OTD) measures. Using Call Report data for 2017–2020, we standardize loan assets, trading assets, and repurchase-agreement assets within year and combine them into the OTD score:

$$\text{OTD Score} = -z(\text{Loan Assets}) + z(\text{Trading Assets}) + z(\text{Repo Assets}). \tag{A1}$$

High-OTD banks hold less of their assets in loans and more in trading and repo positions than Low-OTD banks, and account for most datacenter mortgage debt despite representing a small share of datacenter-lending banks.

We validate the classification using matched DealScan lender-tranche data, restricted to domestic banks (Appendix Table B1). High-OTD banks account for a substantially larger share of matched DealScan deal volume, are less likely to participate via full syndication, and are more likely to participate in club deals, consistent with a more distribution-oriented lending model.

Foreign banks. For most foreign banks we lack usable Call Report data, so we classify their High- versus Low-OTD status using DealScan syndication behavior instead. The DealScan classification uses a composite index of full-syndication share, club-deal share, and deal volume and count, calibrated against domestic banks with an unambiguous Call Report-based classification (Appendix Table B1), and is applied explicitly to Mizuho, Royal Bank of Canada, Sumitomo Mitsui, Société Générale, Natixis, CIBC, MUFG, and Bank of Montreal. Barclays, Deutsche Bank, Toronto Dominion, and Royal Bank of Scotland instead retain a Call Report-based classification, independently confirmed by the same DealScan composite. Eurohypo, with no DealScan coverage for this entity, is classified Low given lack of presence in DealScan. Lehman Brothers Holdings predates our DealScan sample window entirely and is classified as High based on its business model.

The three DealScan behaviors reported in Appendix Table B1, Panel A, are constructed directly from DealScan's own fields. Full syndication and club deal are indicators for DealScan's `DistributionMethod` field equal to "Syndication" or "Club Deal" respectively. Lead arranger is an indicator for DealScan's `PrimaryRole` field equal to Arranger, Bookrunner, Lead arranger, Mandated Lead arranger, Admin agent, or Syndication agent, as opposed to Participant.

Syndicated and CMBS lenders. CMBS conduit and syndicated lenders are not covered by Call Report data and therefore have no directly observed OTD score. In our binary

classification, syndicated lenders are rule-coded High-OTD and CMBS/non-bank lenders are classified Low-OTD, so this absence of a Call Report score does not require imputation for any result we report.

Bank size variables. We also construct predetermined bank-size measures from total assets as of 2021 Q4, using Call Report data for domestic banks and publicly available balance sheet data for foreign institutions, as robustness controls in Appendix Table D4: a continuous $\ln(\text{Assets}_{2021})$ control (Panel A) and an exclusion of the 15 largest lenders (Panel B).

Acquired and renamed banks. Several lenders in our sample were acquired by or merged into other institutions before the end of our sample period. In most cases, including SouthTrust and SunTrust, loan-level debt and OTD classification are both attributed to the lender name recorded at origination, preserving each predecessor's own historical identity rather than mapping it to the acquiring institution's current balance sheet. One exception is Bank One, which is relabeled to its acquirer, JPMorgan Chase, for every loan in our sample regardless of origination date, since Bank One's own Call Report identity is not separately recoverable in our data. One loan in our sample, originated in 1997, is labeled JPMorgan Chase under this convention; excluding this loan does not change any reported result.

Lender clustering. Standard errors are clustered on an integer identifier assigned to each standardized lender name. Unlike the historical acquisitions above, where predecessor identity is preserved, a subsidiary that is currently part of an active parent's ongoing lending business is standardized to the parent's name before classification and clustering (for example, CIT Bank N.A. and CIT Group; German American Capital and Deutsche Bank; BB&T and Truist), so that a single lending relationship is not split across two identities with potentially different OTD scores.

A.4 Cleaning and Harmonization of Property Variables

Duplicates and normalization. We eliminate duplicate property records and retain one observation per property in the main property-level file. Descriptive fields (address, market, submarket, county, and borrower names) are trimmed and normalized (case and punctuation) to ensure consistent matching and reporting across sources.

Rent per square foot. CoStar reports shell rent as `rentsfyr` (a string field). We convert this field into a numeric rent per square foot measure, imputing missing or invalid values using the same county-then-state-then-national median hierarchy used for property taxes below. For summary statistics and regressions, rent is interpreted as shell (base) rent and excludes colocation usage charges and other operating pass-throughs.

Property taxes per square foot. CoStar reports property taxes as `taxespersf` (string) and/or `taxestotal`. We construct a numeric taxes-per-square-foot measure by removing punctuation and currency symbols and converting to a numeric value. When `taxespersf` is not reported, we impute it using the median of observed `taxespersf` among datacenters in the same county, falling back to the state median and then the national median if the county (or state) median is itself unavailable.

Percent leased. CoStar reports percent leased as `percentleased` (string). We harmonize this field into a numeric 0–100 measure by stripping percent symbols and converting to numeric values. Values between 0 and 1 are treated as fractions and scaled to percentages. Values outside the 0–100 range are set to missing.

Building age. CoStar provides building age measured as of 2025 (`age_2025`). We use this variable directly and interpret it as the number of years since construction, measured as of 2025. Observations with implausible values (negative ages or ages exceeding 150 years) are set to missing.

Development status. CoStar reports building status in nine categories: Existing, Under Renovation, Under Construction, Proposed, Final Planning, Abandoned, Demolished, Converted, and Deferred. The regression sample has no missing status. Appendix Table D6,

Panel C, excludes properties coded Abandoned, Proposed, Final Planning, or Under Construction, since these may not yet be, or may no longer be, operational; all other statuses remain in the sample.

A.5 Construction of Property-Level Debt and Leverage Measures

Debt allocation for multi-property loans. Loan origination records may correspond to a loan secured by multiple properties. To construct property-level debt intensity measures, we apportion each loan's balance across collateral properties in proportion to each property's rentable building area (RBA). Formally, if loan ℓ is secured by properties $i \in \mathcal{I}_\ell$ with RBA RBA_i, the debt allocated to property i from loan ℓ is

$$Debt_{i\ell} = Debt_\ell \times \frac{RBA_i}{\sum_{j \in \mathcal{I}_\ell} RBA_j}.$$

Property-level debt is the sum of allocated balances across loans collateralized by that property. RBA-proportional allocation is the economically appropriate way to assign a portfolio loan across its pledged facilities. Absent property-specific balance data, physical scale is the most direct observable basis for splitting a balance that jointly finances multiple properties.

For brevity, the main text refers to allocated property-level debt simply as "debt." "Origination amount" refers specifically to the original CoStar-reported loan amount prior to RBA allocation across collateral properties.

Debt per square foot. Debt per square foot is computed as allocated property debt divided by property RBA. This measure is defined for the subset of properties with positive allocated mortgage balances.

Debt-to-rent. Debt-to-rent ratios are constructed using shell rent per square foot.

Loan-to-value (LTV). LTV is defined as mortgage balance divided by real-estate collateral value. Real-estate collateral value is the maximum of three approaches: government tax assessment from matched parcel records, with a state-median assessed value per square foot

used as a fallback for properties lacking a usable direct assessment; an income-based value capitalizing CoStar shell rent over a finite horizon; and an acre-based land value using the state median dollars-per-acre among properties with matched assessments. Formally,

$$\text{Value}_i \equiv \max\left\{\text{Assess}_i,\ \text{DCF}_i,\ \text{Land}_i\right\},$$

and

$$\text{LTV}_i = \frac{\text{AllocatedDebt}_i}{\text{Value}_i}.$$

The maximum implements the highest-value-use principle. When the property supports more than one valuation basis, its market value is the highest realizable value rather than a lower alternative appraisal. The income-based value DCF_i capitalizes shell rent using a 10-year explicit discounted cash flow with an annual rent growth rate of $g = 10\%$ and a discount rate of $r = 7\%$, plus a terminal value at year 10 based on a lower long-run terminal growth rate of $g_{\text{terminal}} = 2\%$, implying a 5% terminal capitalization rate. These parameters are applied consistently across properties. The non-datacenter placebo sample of Section 6.2 lacks matched parcel records, so Value_i there uses $\max\{\text{DCF}_i, \text{Land}_i\}$, omitting the assessment component.

Origination timing. CoStar reports the loan origination date as `originationdate` (string). We parse this into a daily date when possible and extract the origination year. Years since origination is defined as $2026 - \text{OriginationYear}_i$.

Transformations used in estimation. We use log transformations of continuous outcomes: ln(DebtPSF), ln(DebtToRent), and ln(LTV) as dependent variables, and ln(RBA) as a baseline control. When constructing logged outcomes, we restrict to observations with strictly positive values.

Extreme loan-to-value observations. The maximum LTV in the sample is 73.9, reported in Table 1 (full collateral-value construction in Appendix Table A1). This observation and a small number of similar cases reflect genuine single-property loans to specialized datacenter

operators in which the real-estate collateral value, computed as the maximum of assessment, income capitalization, and land value, is very low relative to the loan balance. These loans include financing to known operators such as EdgeConneX, QTS, and ServerFarm, and represent structures in which lenders appear to underwrite against expected cash flows and platform value rather than in-place collateral. They are not data or allocation errors; the leverage measures are constructed correctly, and the extreme-value checks in Appendix Table C2 confirm that the main results are unaffected by their inclusion.

Excluded observation. One Flexential property is excluded from the sample. CoStar records a $2.0 billion loan from an unidentified private lender on a 33,588 square foot facility in February 2022, implying debt per square foot of $59,545. The only contemporaneous Flexential financing in the public record is an $800 million asset-backed securities offering, a different instrument and amount. Because the lender is unidentified and the implied leverage is inconsistent with available public information, we cannot verify this observation and exclude it to avoid inflating the leverage statistics.

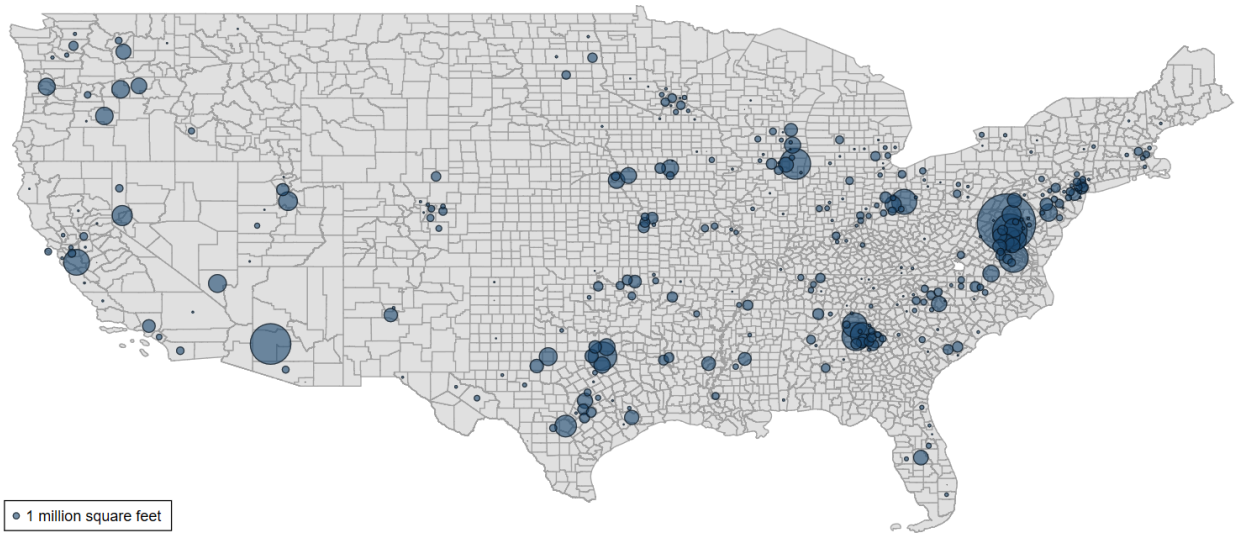


Figure A1. Datacenter Capacity by County

CoStar inventory as of May 2026. Total rentable building area across existing and under-construction properties; marker size proportional to county capacity.

Table A1. Debt Exceeds Real-Estate Collateral Value

Variable	N	Mean	SD	Min	Max
Panel A: Collateral and debt per square foot (USD)					
Tax assessment value per square foot	381	585.4	1,728.4	11.7	29,574.2
Discounted cash flow value per square foot	610	501.0	328.9	80.3	4,309.6
Acre-based land value per square foot	668	658.3	2,174.5	1.5	45,008.9
Real-estate collateral value per square foot (final basis)	668	1,003.5	2,424.6	41.3	45,008.9
Debt per square foot	668	1,436.0	3,463.8	0.0	31,866.6
Panel B: Implied loan-to-value ratios					
Tax assessment basis	381	18.5	61.3	0.0	557.9
Discounted cash flow basis	610	3.7	8.8	0.0	84.7
Acre-based land value basis	668	13.4	54.4	0.0	890.5
Final collateral basis	668	2.8	7.4	0.0	73.9

Note: Real-estate collateral value construction follows Appendix A.5. The tax assessment row reports direct parcel-matched assessments only ($N = 381$); the state-median assessed-value-per-square-foot fallback used for properties without a direct match is not separately broken out but is reflected in the final-basis row ($N = 668$). The final-basis LTV distribution (median, SD, min, max) is reported in Table 1, Panel B; 36.5% of properties carry debt exceeding their collateral value.

Appendix B: Lender Classification

Table B1. DealScan Validation

Panel A: Deal-level syndication behavior			
	Full Syndication	Club Deal	Lead Arranger
High-OTD bank	−0.040***	0.037***	0.190*
	(0.009)	(0.008)	(0.092)
Log deal size	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
R^2	0.035	0.037	0.061
Observations	3,954	3,954	3,954
Panel B: Bank-level scale			
	High-OTD	Low-OTD	Difference
Mean DealScan volume per bank ($bn)	636.5	88.7	547.8***

Note: Sample restricted to banks with a matched OTD score, using DealScan lender-rows matched to the CoStar lender universe. Panel A reports linear probability models of the form

$$\text{Behavior}_i = \beta\,\text{HighOTD}_i + \theta\,\ln(\text{DealSize}_i) + \delta_t + \epsilon_i,$$

with year fixed effects δ_t. DealScan includes syndicated facilities and is not restricted to mortgage debt. The omitted category is Low-OTD banks. Standard errors clustered at the lender level are reported in parentheses. Panel B reports mean bank-level matched DealScan deal volume; the difference uses an unequal-variance test. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table B2. Borrower Classification

Undiversified Datacenter		Big Tech	Government	Other
QTS	BlackChamber Group	Amazon (AWS)	Scott County Board of Supervisors	Blackstone REIT
Digital Realty Trust, Inc.	Centra Digital	Microsoft		Landmark Dividend
CyrusOne	Equinix	Google		Brookfield
EdgeConneX	Redwolf DCD KC II LLC	Meta		GI Partners
Mapletree Industrial Trust	Rowan Digital Infrastructure			IFM Investors
STACK Infrastructure	ScaleMatrix			Menlo Equities
DataBank	Sentinel Data Centers, LLC			Igneo Infrastructure Partners
Prime Data Centers	TRG Datacenters			Core Scientific, Inc.
Vantage Data Centers	ValorC3 Data Centers			StratCap
Aligned Data Centers	AiNET			Aphorio Carter
Switch, Inc.	CDS Global Cloud			Chirisa
Stream U.S. Data Centers, L.L.C.	Cirrascale Cloud Services			AREP
H5 Data Centers	Coloswitch Llc			Oppidan
NTT Data Group Corporation	DartPoints			fifteenfortyseven Critical Systems Realty
T5 Data Centers	Data Center Leasing, Inc.			5Rivers CRE
CSquare	Datacate Inc			Blue Owl Capital
Cologix	Direct LTx			Mark Krohn
Edged	Eastern Iowa Data Center Llc			Priseda
Flexential	Evocative			Stonepeak Partners LLC
CoreSite	Five 9s Digital			Various other single-property owners*
Digital Core REIT	Gennex Group			
Iron Mountain Incorporated	Involta LLC			
Serverfarm	Lifeline Data Centers			
TierPoint	Managedway			
US Signal	Markley Group			
EdgeCore Digital Infrastructure	NC Datacenter, LP			
Novva Data Centers	NV RNO 1 LLC			
Compass Datacenters	Next Tier HD			
DC BLOX	PowerHouse Data Centers			
Netrality Data Centers	Rackspace			
CloudHQ	Red Wolf DCD Properties LLC			
DigiCo Infrastructure REIT	Scale Datacenter			
Element Critical	Skybox Datacenters			
Lightedge	Yondr			
Sabey Data Centers	eVolve Data Center Solutions			
5C Data Centers	phoenixNAP, LLC			
Ark Data Centers				

Note: Owners within each classification are sorted by number of matched properties in the sample, descending. *114 further owners with a single property each.

Table B3. Lender Classification

High-OTD lenders	Low-OTD lenders	Non-bank lenders	Government lenders
Wilmington Trust	Truist Bank	Silver Point Finance	CoBank ACB
Goldman Sachs	Citizens Bank	Integrated Housing Solutions	State of Ohio
Toronto Dominion Bank	U.S. Bank	Argentic Real Estate Finance	Farm Credit Mid-America FLCA
Wells Fargo	Barclays	Column Financial	Farm Credit of the Virginias ACA
Citibank	KeyBank	Roserock Capital Fund I	Badgerland Financial FLCA
Bank of America	First Bank	Glas Americas	Compeer Financial
JPMorgan Chase	SunTrust Bank	Principal Life Insurance	Small Business Administration
Sumitomo Mitsui Banking	Sunflower Bank	Fidelity Guaranty Life Insurance	County of Washington
Morgan Stanley	First Citizens Bank & Trust	Alter Domus US	Iowa Finance Authority
Computershare	Comerica Bank	ING Capital	
Fifth Third Bank	Huntington National Bank	New York Life Insurance	
Société Générale	Arizona Business Bank	Pacific Life Insurance	
Deutsche Bank	Hancock Whitney Bank	American General Life Insurance	
CIT Group	Valley National Bank	AIG Asset Management US	
UMB Bank	Pinnacle Bank	Acore Capital Mortgage	
Royal Bank of Canada	Bankers Trust	Real Estate Capital Services	
CIBC	Eagle Bank	Various other non-bank lenders†	
MUFG Bank	Webster Bank		
Natixis New York Branch	Pacific Western Bank		
Compass Bank	PNC Bank		
BNP Paribas	UBS AG		
Bank of Montreal	Capital One		
Mizuho Bank	First Citizens Bank		
Royal Bank of Scotland	Eurohypo AG New York Branch		
Lehman Brothers Holdings	Various other Low-OTD lenders**		
Merrill Lynch			
Horizon Bank			
BOK Financial			
Commerce Bank			

Note: Lenders within each classification are sorted by number of matched properties in the sample, descending. **91 further Low-OTD lenders. †58 further non-bank lenders.

Appendix C: Leverage and Borrower Selection

Table C1. Preboom Equality Tests for Event Study Coefficients

Outcome	Joint F	Joint p	-3 p	-2 p	Obs.	Lender clusters
Debt per Square Foot	0.121	0.948	0.554	0.636	668	228
Debt to Rent Ratio	0.165	0.920	0.592	0.488	668	228
Loan to Value Ratio	0.309	0.819	0.443	0.379	668	228
Undiv. DC Count	1.605	0.189	0.638	1.000	5,928	228
Loan Amount	0.167	0.919	0.512	0.533	668	228

Note: Each row reports the joint test of the null that the High-OTD gaps in all pre-shock windows equal the period -1 (immediately pre-shock) reference-window gap, using the same specification as the corresponding event-study figure: Figure 3 for the leverage and loan-amount rows, Figure 4 for the Undiv. DC Count row. The next two columns report the individual tests for periods -3 and -2. Standard errors are clustered at the lender level.

Table C2. Robust to Outliers

	ln(DebtPSF)	ln(DebtToRent)	ln(LTV)	Undiv. DC Count
Winsorized (p1/p99)				
High-OTD × Post-Jul'21	1.548***	1.472***	1.464***	1.179***
	(0.405)	(0.407)	(0.433)	(0.425)
Controls and FE	Yes	Yes	Yes	Yes
Observations	668	668	668	10,556
Trimmed (top 1%)				
High-OTD × Post-Jul'21	1.663***	1.514***	1.519***	1.165***
	(0.441)	(0.420)	(0.457)	(0.419)
Controls and FE	Yes	Yes	Yes	Yes
Observations	661	662	662	10,452

Note: This table reports estimates from the baseline DiD specification (Equation 2) under two adjustments for extreme observations: winsorizing leverage at the 1st and 99th percentiles, or dropping observations above the 99th percentile. The Undiv. DC Count column reports the Table 4 specification. The three leverage columns include the baseline controls and fixed effects in Equation 2; the Undiv. DC Count column instead uses Table 4's lender-window-state Poisson design and controls. Standard errors clustered at the lender level are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Appendix D: Robustness and Alternative Explanations

Table D1. High vs. Low-OTD: Matching Estimates

	ln(DebtPSF)		ln(DebtToRent)		ln(LTV)	
	Prop. Score	Exact	Prop. Score	Exact	Prop. Score	Exact
High-OTD × Post-Jul'21	1.435***	1.174***	1.262***	1.133***	1.391***	1.169**
	(0.432)	(0.434)	(0.420)	(0.427)	(0.385)	(0.449)
Controls and FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	462	502	462	502	462	502

Note: This table reports estimates from the baseline DiD specification (Equation 2), weighted by propensity-score (Prop. Score columns) or exact-matching (Exact columns) weights on state and property-size bins; the sample is restricted to lenders with a defined High- or Low-OTD classification, excluding lenders without a comparable counterfactual group, and matching methodology is described in Section 6.1. Controls and fixed effects are as in Table 3. Standard errors clustered at the lender level are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table D2. Alternative Explanations: Borrower Fixed Effects

	ln(DebtPSF)	ln(DebtToRent)	ln(LTV)
High-OTD × Post-Jul'21	1.673***	1.406**	1.592**
	(0.606)	(0.618)	(0.627)
Borrower FE	Yes	Yes	Yes
Observations	668	668	668

Note: This table adds borrower fixed effects to the baseline controls and fixed effects of Equation 2, identifying the High-OTD coefficient from each borrower's own pre-to-post change; identification is discussed in Section 6.1. Standard errors clustered at the lender level are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table D3. Non-Datacenter Commercial Real Estate Placebo

	ln(DebtPSF)	ln(DebtToRent)	ln(LTV)
Panel A: Leverage outcomes (non-datacenter industrial commercial real estate)			
High-OTD bank	−1.875***	−1.755***	−1.624***
	(0.528)	(0.503)	(0.417)
Controls and FE	Yes	Yes	Yes
Observations	886	883	886
Panel B: Deal scale			
	ln(Loan Amount)	ln(RBA)	
Datacenter	1.504***	0.229**	
	(0.266)	(0.100)	
Non-Datacenter	−0.927*	0.214	
	(0.533)	(0.262)	
Controls and FE	Yes	Yes	
Observations	668 / 94	668 / 924	

Note: This table reports the pooled cross-sectional analogue of the baseline DiD specification, omitting the High-OTD × Post interaction since this non-datacenter placebo sample (Appendix A.2) has no datacenter-boom timing to benchmark against. The sample is restricted to bank lenders. Both panels estimate

$$Y_{it} = \beta_H \, \text{HighOTD}_i + \mathbf{X}_{it}'\theta + \gamma_s + \delta_t + \varepsilon_{it},$$

with HighOTD_i, $\mathbf{X}_{it}$, γ_s, and δ_t defined as in Equation 2, except Panel B drops RBA as a control, since loan amount and RBA are themselves the outcomes there. Because Panel B is a pooled cross-sectional comparison, its datacenter RBA coefficient is not directly comparable to the DiD estimate of the post-boom change in RBA reported in Table 7. Standard errors clustered at the lender level are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table D4. Alternative Explanations: Bank Size

	ln(DebtPSF)	ln(DebtToRent)	ln(LTV)	Undiv. DC Count
Panel A: Predetermined bank-assets control, bank lenders only				
High-OTD × Post-Jul'21	1.713***	1.661***	1.744***	1.525***
	(0.414)	(0.433)	(0.461)	(0.546)
$\ln(\text{Assets}_{2021})$	0.071	0.070	0.072	−0.014
	(0.041)	(0.043)	(0.043)	(0.043)
Observations	488	488	488	6,095
Panel B: Excluding the 15 largest lenders				
High-OTD × Post-Jul'21	1.344***	1.404***	1.301***	1.448**
	(0.455)	(0.463)	(0.493)	(0.606)
Observations	494	494	494	8,632
Controls and FE	Yes	Yes	Yes	Yes

Note: This table reports estimates from the baseline DiD specification (Equation 2) under two alternative treatments of bank size, distinguishing OTD incentives from generic large-bank behavior as a driver of the leverage premium. Panel A includes bank lenders only; Panel B uses the full sample. Controls and fixed effects for the three leverage columns are as in Table 3; the Undiv. DC Count column instead uses Table 4's lender-window-state Poisson design and controls. Standard errors clustered at the lender level are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table D5. Classification Robustness

	ln(DebtPSF)	ln(DebtToRent)	ln(LTV)	Undiv. DC Count
Panel A: 2022–2025 average classification, same overrides				
High-OTD × Post-Jul'21	1.533***	1.473***	1.415***	1.218***
	(0.408)	(0.412)	(0.437)	(0.426)
Observations	668	668	668	10,556
Panel B: Excluding Syndicated-type lenders				
High-OTD × Post-Jul'21	1.650***	1.589***	1.586***	1.032**
	(0.454)	(0.456)	(0.484)	(0.410)
Observations	576	576	576	9,724
Panel C: Excluding foreign banks				
High-OTD × Post-Jul'21	1.457***	1.406***	1.398***	1.078**
	(0.424)	(0.436)	(0.453)	(0.488)
Observations	584	584	584	9,438
Controls and FE	Yes	Yes	Yes	Yes

Note: This table re-estimates the baseline DiD specification (Equation 2) and Table 4's Poisson borrower-selection design under three alternative treatments of the High-OTD classification. Panel A replaces only the classification window: every Bank-type lender is reclassified using the 2022–2025 Call Report decile rule in place of the primary 2017–2020 window, applying the identical manual overrides described in Appendix A.3 (foreign banks classified from DealScan behavior, and a small number of domestic banks reclassified on other evidence); banks without direct 2022–2025 Call Report coverage default to Low-OTD absent an override, mirroring the primary classification's own default rule. This isolates the window choice as the sole difference from the primary classification. Panel B excludes Syndicated-type lenders, which are rule-coded High-OTD rather than scored from Call Report data. Panel C excludes foreign banks, classified using DealScan syndication behavior rather than Call Report data (Appendix A.3). Standard errors clustered at the lender level are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table D6. Alternative Debt Allocation and a Sample Restriction

	ln(DebtPSF)	ln(DebtToRent)	ln(LTV)	Undiv. DC Count
Panel A: Single-property loans				
High-OTD × Post-Jul'21	1.484***	1.459***	1.420***	0.966**
	(0.426)	(0.420)	(0.460)	(0.406)
Observations	465	465	465	9,490
Panel B: Equal-split allocation				
High-OTD × Post-Jul'21	1.738***	1.663***	1.656***	1.179***
	(0.423)	(0.427)	(0.450)	(0.425)
Observations	668	668	668	10,556
Panel C: Excluding abandoned, proposed, final-planning, and under-construction properties				
High-OTD × Post-Jul'21	1.456***	1.377***	1.466***	1.559***
	(0.404)	(0.416)	(0.437)	(0.451)
Observations	554	554	554	8,970
Controls and FE	Yes	Yes	Yes	Yes

Note: This table reports estimates from the baseline DiD specification (Equation 2) under two alternative treatments of debt allocation and a sample restriction. Panel A restricts to single-property loans, eliminating any role for the allocation rule. Panel B allocates loan balances equally, rather than proportionally by RBA, across collateral properties, addressing the concern that the leverage premium is a mechanical artifact of the allocation rule. Panel C drops properties classified by CoStar as abandoned, proposed, final planning, or under construction; all other development statuses, including existing, under renovation, converted, deferred, and demolished, remain in the sample. Controls and fixed effects for the three leverage columns are as in Table 3; the Undiv. DC Count column instead uses Table 4's lender-window-state Poisson design and controls. Standard errors clustered at the lender level are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table D7. Comparison to Purnanandam (2011)

	Purnanandam (2011)	This paper
Shock	2007 collapse of the secondary mortgage market (supply-side)	AI-driven datacenter boom (demand-side)
Distribution measure	OTD share of originations sold or securitized	Bank balance-sheet OTD score (Equation A1)
Identification	High- vs. low-OTD banks, pre/post-2007 comparison (DiD)	High- vs. low-OTD banks, before vs. from July 2021 onward (DiD)
Outcome	Loan performance (delinquency, default)	Loan structure at origination (leverage, borrower selection)
Matching/controls	Matched banks and borrowers	Coarsened exact matching, IPW
Instrument	N/A	2020 loan-to-assets ratio
Placebo	N/A	Pre-boom (2015–2019); non-datacenter industrial commercial real estate
Pricing dispersion test	Spread dispersion across OTD groups (HMDA)	Not feasible: datacenter mortgage pricing not systematically disclosed

Appendix E: Instrumental-Variables Analysis

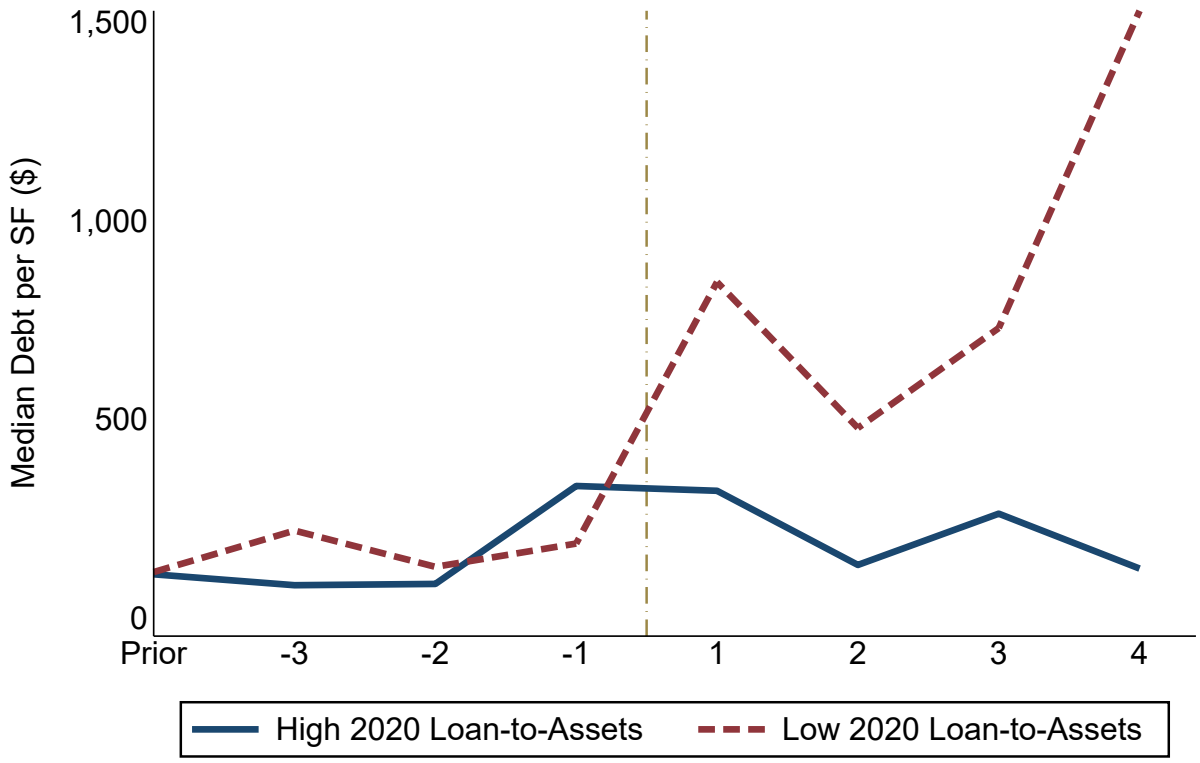


Figure E1. Leverage Levels by Instrument Split

Median debt per square foot by lender group, anchored at the July 2021 shock. Post-shock periods match Figure 3's four 15-month bins; pre-shock periods use four 3-year bins given this instrument-matched sample's thinner coverage. Lenders are split at the lender-level median of the predetermined 2020 loan-to-assets ratio, the excluded instrument.

Table E1. Persistence of OTD Classification

Later year	2021	2022	2023	2024	2025
Instrument persistence (cross-bank correlations)					
Loan-to-assets ratio	0.966	0.925	0.921	0.916	0.873
Banks	103	103	103	103	103

Note: Cross-bank correlations between the 2020 loan-to-assets ratio and its values in subsequent years, computed once per matched bank.

Table E2. Exclusion Restriction Diagnostics

	Coefficient on instrument	SE	p-value
Panel A: Property Balance Test			
ln(RBA)	<0.001	<0.001	0.800
Rent per square foot	2.454	1.669	0.145
Taxes per square foot	<0.001	<0.001	0.565
Panel B: Pre-Boom Market Share			
ln(2015–2019 datacenter loan volume)	−4.487	1.762	0.016**
2015–2019 datacenter market share (fraction)	−0.085	0.068	0.223
Panel C: Bank-Level Balance Test			
Return on assets	−0.003	0.014	0.819
Net interest margin	0.025	0.003	0.000***
Loan-loss provision ratio	0.001	0.002	0.349

Note: The table reports diagnostic OLS estimates for bank lenders. Each panel regresses the indicated variable on the instrument Z_b, following the structure of Equation 3 with the diagnostic variable in place of HighOTD_b; standard errors are clustered at the lender level in Panel A and heteroskedasticity-robust in Panels B and C. Panel C restricts to banks that appear as lenders in the estimation sample, using 2020 Q4 Call Report data; net interest margin scales mechanically with loan volume and is not a clean quality signal, so return on assets is the relevant test. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table E3. Instrumental-Variable Estimates with Funding Controls

Funding control added	ln(DebtPSF)	ln(DebtToRent)	ln(LTV)	Obs.
+ Brokered/assets	1.351***	1.472***	1.295***	475
	(0.413)	(0.408)	(0.393)	
+ Int. exp./dep.	1.692***	1.816***	1.674***	467
	(0.546)	(0.535)	(0.507)	
+ Int. exp./liab.	1.801***	1.922***	1.768***	475
	(0.554)	(0.544)	(0.519)	
+ Deposits/assets	1.309***	1.431***	1.246***	475
	(0.401)	(0.392)	(0.360)	

Note: The table reports 2SLS estimates for bank lenders only. Each row adds the indicated 2020 funding control, one at a time, to Equation 4, building on Table 6. Entries are coefficients on High-OTD bank; all funding controls are standardized. Standard errors clustered at the lender level are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table E4. Pre-Boom Placebo and Vintage Robustness

	ln(DebtPSF)	ln(DebtToRent)	ln(LTV)	KP-F
Panel A: Pre-Boom Placebo (2015–2019)				
High-OTD bank	0.021	0.369	0.511	4.34
	(0.954)	(0.970)	(0.980)	
Controls and FE	Yes	Yes	Yes	
Observations	56	56	56	
Panel B: Vintage Robustness (Alternate Instrument Years)				
Instrument vintage	ln(DebtPSF)	ln(DebtToRent)	ln(LTV)	KP-F
2015–2017 average	1.301***	1.424***	1.235***	36.41
	(0.447)	(0.431)	(0.407)	
2018–2019 average	1.245***	1.364***	1.157***	23.88
	(0.414)	(0.403)	(0.374)	

Note: Both panels report 2SLS estimates for bank lenders only using Equations 3–4. Panel A applies the 2020 loan-to-assets instrument to the 2015–2019 pre-boom sample; baseline results, estimated on the full sample, are reported in Table 6. Panel B re-estimates that same baseline specification using earlier-vintage instruments in place of the 2020 value, on the full outcome sample. Standard errors are clustered at the lender level and reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table E5. IV Estimates: Borrower Selection

	Undiv. DC Count
Bank-sample (uninstrumented)	1.601***
	(0.580)
2SLS: Cross-sectional IV	0.081***
	(0.026)
Controls and FE	Yes
Observations	6,578
KP-F	35.65

Note: The sample is restricted to bank lenders matched to Call Report data. 2SLS estimates extend Equations 3–4 to the extensive-margin borrower-selection outcome (Baseline: Table 4 specification). High-OTD status is instrumented by the 2020 loan-to-assets ratio. The 2SLS coefficient is linear and not comparable in scale to Baseline, since two-stage least squares does not support a nonlinear count model. Standard errors are clustered at the lender level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table E6. IV Estimates: Borrower-Selection Robustness

	Undiv. DC Count	KP-F	Obs.
Panel A: Funding Controls (2SLS, one added at a time)			
+ Brokered/assets	0.082***	72.69	6,578
	(0.024)		
+ Int. exp./dep.	0.093***	24.95	6,448
	(0.027)		
+ Int. exp./liab.	0.085***	33.40	6,578
	(0.025)		
+ Deposits/assets	0.087***	26.98	6,578
	(0.024)		
Panel B: Pre-Boom Placebo (2015–2019)			
High-OTD bank	0.006	35.75	5,566
	(0.006)		
Panel C: Vintage Robustness (Alternate Instrument Years)			
2015–2017 average	0.076***	58.16	6,578
	(0.023)		
2018–2019 average	0.079***	33.07	6,578
	(0.026)		

Note: This table extends the leverage funding-control, pre-boom placebo, and instrument-vintage checks (Appendix Tables E3 and E4) to the extensive-margin borrower-selection outcome, mirroring Table E5's specification. The sample is restricted to bank lenders matched to Call Report data. Panel A adds the indicated 2020 funding control, one at a time, to the cross-sectional 2SLS specification; all funding controls are standardized. Panel B applies the 2020 loan-to-assets instrument to the 2015–2019 pre-boom sample. Panel C uses earlier-vintage instruments in place of the 2020 value, on the full outcome sample. Standard errors clustered at the lender level are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Appendix F: Mechanism: Loan Size

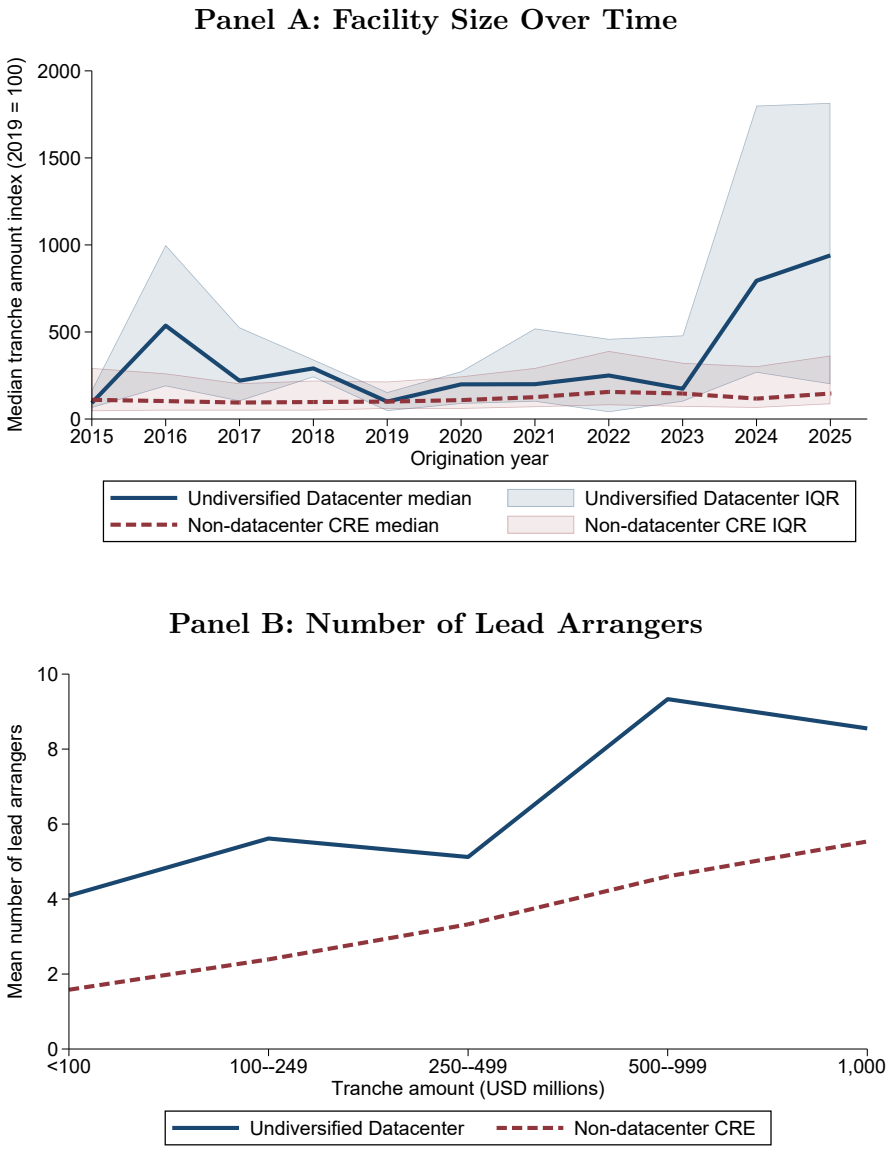


Figure F1. DealScan Confirmation: Loan Growth and Syndicate Structure

Undiversified Datacenter facilities (all loan purposes) versus non-datacenter industrial CRE facilities (real-estate/construction purpose only). Panel A: median facility amount and interquartile range, indexed to 2019, through 2025. Panel B: mean lead arrangers by facility-size category, 2015–March 2026. Sample covers syndicated facilities, not only mortgage debt.

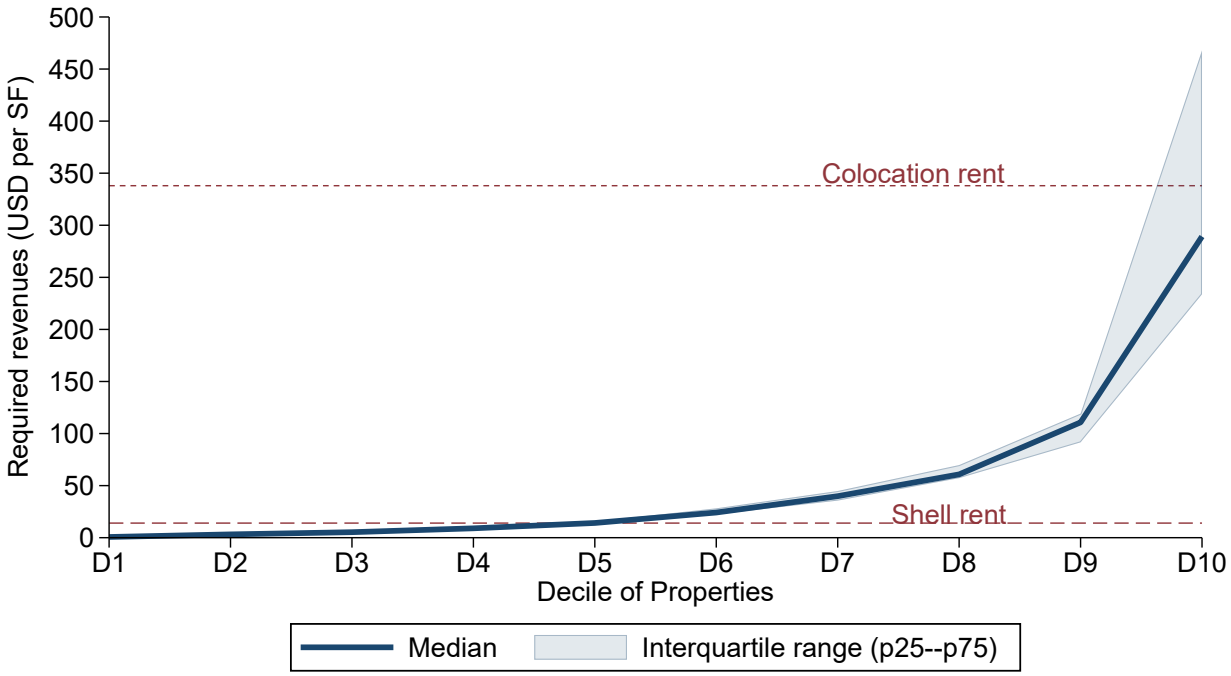


Figure F2. Required Revenue by Debt Burden

CoStar datacenter properties with mortgage debt as of May 2026 are ranked into deciles by debt per square foot. Required revenue is annual interest expense per square foot plus property taxes per square foot. The solid line reports the median within each decile; the shaded band reports the interquartile range. Horizontal dashed lines indicate median shell rent ($14 per square foot, CoStar) and benchmark colocation rent ($338 per square foot, Digital Realty 3Q25 earnings release, October 23, 2025, GAAP base rent per square foot on new leases signed, Americas 0–1 MW retail colocation segment).

Table F1. Fee Structure in Datacenter Credit Facilities

Recipient	Fee	Contractual basis
Phoenix Data Center Acquisitions LLC (2024)		
Lead arrangers	Upfront fees and OID	Confidential fee letter
Revolving lenders	Commitment fee	0.75% p.a. of undrawn commitment (flat)
L/C participant lenders	Participation fee	4.25% p.a. of avg. daily L/C exposure, less fronting fee
L/C issuing bank	Fronting fee	0.25% p.a. of outstanding L/C amount
Administrative agent	Agency fee	Separately negotiated
CoreSite, L.P. (2019)		
Lead arrangers	Upfront fees	Confidential fee letter
Revolving/term lenders	Unused facility fee	0.25% p.a. of undrawn commitment if drawn $\leq$ 50%, 0.15% p.a. if $>$ 50% (steps down further at Investment Grade Pricing Date)
Administrative agent	Agency fee	Separately negotiated

Note: Fee provisions from publicly filed credit agreements (CoreSite, L.P., 2019; Phoenix Data Center Acquisitions LLC, 2024). Upfront arranger fees and administrative-agent compensation are set by confidential fee letters not filed with the SEC.