Data Center Mortgages and Originate-to-Distribute Lenders*

Pedro Gete†, Amparo Mercader‡

October 2026

Abstract

We build a novel property-level database of U.S. data center mortgages, an opaque financing channel of the AI buildout, linked to lender balance sheets. A few distribution-oriented (High-OTD) lenders originate most of this credit. After mid-2021, the ratio of High-OTD to other-lender debt relative to observable real estate rises five- to sixfold. The increase reflects larger loans on properties of similar size. High-OTD lenders also lend more to specialized data center borrowers, and the increase in debt is larger for these borrowers. More of this debt therefore depends on data center demand.

*We thank Aleksandar Andonov, Aymeric Bellon, Olivier Darmouni, Juan Dubra, Richard Evans, Manuel Garcia, Pablo M. Garcia, Carlos Garriga, Marco Giacoletti, Juan Pedro Gómez, Robert Kurtzman, Enrique Martinez, Carlos Ramirez, Athena Tsouderou, Victoria Vanasco, Stijn Van Nieuwerburgh, Isabel de Valenzuela, and participants at the 2026 European Finance Association, the 2026 Dallas Fed Conference, the University of Alicante, and the 2026 Spanish Finance Conference for helpful comments.

This paper was partially funded by MICIU/AEI/10.13039/501100011033/FEDER, UE, Grant No. PID2024-161318NB-I00. We have read the Journal of Finance disclosure policy and have no conflicts of interest to disclose.

†IE University. Email: pedro.gete@ie.edu.

‡Georgetown University. Email: mercada@georgetown.edu

The AI-driven data center buildout is one of the largest capital investment expansions in recent history (Wachter and Wachter, 2026). Existing work studies the corporate and sponsor financing of the buildout (Van Nieuwerburgh, 2026). We study the mortgages secured by the data centers themselves, and ask who supplies this credit, to whom, and how lending changed as the boom began. Data centers are a new asset class that has not yet faced a downturn.

Gross data center mortgage originations since January 2015 total $158 billion, yet little is known about this market. Data center mortgages are spread across property records, local lenders, and private loan documents, with no central disclosure like the one for residential mortgages or corporate bonds. We build the first property-level dataset that links U.S. data center mortgages to lender balance sheets. To study how credit depends on lenders' business models, we classify banks by their originate-to-distribute (OTD) orientation. Following Purnanandam (2011) and Adrian and Shin (2010), we construct OTD scores from predetermined bank balance sheets and classify the most distribution-oriented banks as High-OTD.

We find that a few lenders originate most of this credit and a few borrowers account for most of it. High-OTD lenders originate 87.9% of data center mortgage credit since 2015, including every loan of $1 billion or more. Most of the credit goes to companies whose main business is owning and operating data centers, which we call Undiversified Data Center borrowers, and since the boom began the five largest borrowers account for almost half of it. Hyperscalers, the large technology firms that operate cloud and AI computing at scale, own about a quarter of data center space but account for almost none of the mortgage debt. They have described third-party capacity as a bridge while their own capacity comes online.[1] Demand for leased data center space may therefore be temporary. A drop in that demand would most directly impact the revenues of Undiversified Data Center borrowers.

In a difference-in-differences design, we compare loans by High-OTD lenders and other

[1] Alphabet's chief financial officer stated that "[w]e plan to expand the use of third-party capacity in Q3 as a bridging strategy while we build out more internal capacity." Alphabet Q2 2026 Earnings Call, July 22, 2026.

lenders before and after mid-2021, when data center construction rises sharply. The ratio of High-OTD to other-lender loan amounts per property rises close to fivefold, with no statistically detectable increase in the size of the properties they finance. The ratio of their debt per square foot, debt-to-rent, and loan-to-value rises five- to sixfold. These loans go mainly to Undiversified Data Center borrowers, and the increase in debt relative to real estate at High-OTD lenders is larger for them than for other borrowers. More debt therefore rests on the same real estate. More of it is owed by borrowers whose revenues depend on data center demand.

The results hold when matching on property characteristics, when controlling for technology, within the same borrower and the same lender, and for borrowers that already borrowed from any lender before the boom. The increase is specific to data centers. It does not extend to warehouses, the closest commercial property, even when we compare the same lender's data center and warehouse loans in the same year. The results are also robust to controlling for bank size, excluding the largest banks, alternative OTD classifications, cutoff dates, and outcome definitions, single-property loans, and excluding refinancings.

We contribute to work on financial intermediation and credit allocation, and in particular on the OTD model. Banks choose between holding loans on their balance sheets and selling them, and this choice shapes their credit supply (Buchak et al., 2024). We document how loan size and borrower allocation differ with lenders' distribution orientation when a new asset class expands. We consider three explanations from this literature.

First, lenders expecting to distribute rather than retain credit have weaker incentives to screen and monitor borrowers (DeMarzo and Duffie, 1999; Gorton and Pennacchi, 1995). Securitization is linked to weaker screening and looser credit supply in residential mortgages before the 2008 financial crisis (Keys et al., 2010; Mian and Sufi, 2009; Purnanandam, 2011). In corporate lending, loans sold to investors perform worse (Bord and Santos, 2015) and lead arrangers monitor more when they retain a larger share (Sufi, 2007), although securitized syndicated loans perform no worse than comparable unsecuritized loans (Benmelech,

Dlugosz, and Ivashina, 2012). In commercial real estate, CMBS loans from conduit lenders perform worse than loans from balance-sheet lenders (Black et al., 2012).

Second, distribution can expand lending capacity by freeing balance-sheet space (Irani et al., 2021; Loutskina, 2011) and by lowering funding costs when securitized loans command higher prices (Gete and Reher, 2021). In commercial real estate, larger loans are much more likely to be securitized (Ghent and Valkanov, 2016). Third, lenders specialize in industries where they have expertise (Blickle, Parlatore, and Saunders, 2026), and distribution-oriented lenders may have data center–specific expertise in arranging large financings. Expertise can also lead lenders to finance different borrowers or projects, based on characteristics we do not observe. The combination of higher leverage and greater exposure to sector-dependent borrowers is consistent with weaker underwriting incentives, with greater financing capacity and investor demand, and with specialization and borrower or project selection. Section VI discusses these explanations.

We also contribute to the growing literature on the economics of the AI buildout, which studies its financing architecture (Van Nieuwerburgh, 2026), state tax incentives for data centers (Gargano and Giacoletti, 2025), the response of renewable-energy investment (Heiss et al., 2026), electricity-market spillovers (Darmouni et al., 2026), and the local economic effects of data center entry (Yue and Zeng, 2026). We document how property-level credit for the buildout is supplied and allocated across lenders.

The rest of the paper proceeds as follows. Section I describes the data and Section II the data center mortgage market. Section III presents the empirical strategy and Section IV the results. Section V examines alternative explanations, Section VI discusses interpretation and implications, and Section VII concludes.

I. Data

A. Data Sources

We construct a property-loan-level dataset of U.S. data centers by integrating commercial real estate records, loan origination data, and lender balance sheet information. The core of the dataset comes from *CoStar*, which provides a national inventory of U.S. data center properties[2] as of September 2026 (2,933 assets). For each property, we observe rentable building area (RBA), the building's leasable floor space, 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. From this universe, we identify 854 properties with usable data center mortgage amounts (representing 29.1%). These properties form the mortgaged sample used in the analysis.

Mortgages finance only part of the data center market. Hyperscalers own much of the space but barely borrow through property-level mortgages, and may finance it through corporate bonds and equity instead. Instead, we find that mortgage credit goes mainly to Undiversified Data Center borrowers.[3] Our results describe how lenders finance the part of the buildout that relies on property-level debt.

We link the mortgaged properties to their lenders' quarterly *Call Report* data and to *DealScan* records of syndicated facilities from 2015 through March 2026. Some specifications use different samples because of data availability or design, as noted where each arises. Table I reports summary statistics for the main variables. Appendix A.1 details data sources and sample construction, and Appendix A.2 describes how the estimation sample varies across specifications.

[2]Appendix Figure A1 maps capacity by county.

[3]Appendix Table A.I compares properties with and without an observed mortgage record. Mortgaged properties are smaller and older, more often existing buildings, and more often owned by Undiversified Data Center borrowers.

B. Lender 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 that originates credit to distribute it. We rank all 5,942 U.S. banks with *Call Report* scores by their average score over 2017 to 2020 and classify those in the top two deciles as *High-OTD*. This window predates the July 2021 boom, so the classification could not have been affected by it. Lenders in our sample that file *Call Reports* take the classification of their own score. Foreign banks, which do not file *Call Reports*, are classified by their syndicated lending behavior in *DealScan*, and loans recorded by a securitization trustee or loan agent are classified as High-OTD. Non-bank and government lenders form separate groups. Appendix A.3 describes these steps, and Appendix Table B.I lists all High-OTD and government lenders and the largest Low-OTD and non-bank lenders by number of matched properties. Section V shows the results are robust to an alternative classification window.

Table II, Panel C, shows that High-OTD banks hold less of their assets in loans and more in trading assets and repurchase agreements than Low-OTD banks. In syndicated lending, High-OTD banks more often act as lead arrangers and Low-OTD banks more often participate (Appendix Table B.II). The classification measures a lender's business model, not whether a given data center mortgage is distributed. Appendix Table B.IV documents observed financing structures, but we do not observe subsequent loan sales or the originating lender's retained share.

C. Outcome Measures

We measure leverage by the loan amount per property and by debt relative to the financed real estate. We measure borrower composition by the type of borrower that receives the loan.

Loan amount per property. We construct property-level loan amounts from loan origination records. We first identify whether a loan is secured by more than one property. For these multi-property loans, we use documented property allocations where available and otherwise allocate using RBA or documented collateral shares.

Debt relative to observable real estate. Debt per square foot, debt-to-rent, and loan-to-value benchmark the loan amount per property against three alternative denominators, the size, rent, and value of the financed property. Allocating portfolio loans across their properties makes the measures comparable across single-asset and portfolio loans. 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, which approximates the property's highest-value use and yields the most conservative loan-to-value ratio.

Unlike traditional commercial real estate, where mortgage debt finances land and the building structure, data center loans can also finance specialized equipment such as power, cooling, and network infrastructure. Because specialized equipment, and some sources of repayment such as the operator's service contracts, are not separately observed or verified at the property level in our data, and the value of the equipment is highly uncertain over the medium term, our measures relate debt to observable real-estate collateral and exclude the technology layered on top of it. They therefore capture exposure relative to the building and land under a conservative scenario in which unobserved equipment provides no additional collateral value. The data center supply chain runs from the building's owner to an operator

that may install GPU equipment and sell AI compute services to its customers.[4] The real-estate lease generates shell real-estate rent, which we use in our debt measures.

Borrower allocation. Borrower allocation measures which borrowers receive credit. It counts every new loan and refinancing in the properties' loan histories. We classify borrowers according to the diversification of their revenue base. *Undiversified Data Center* borrowers are firms whose primary business is owning and operating data center real estate. Firms for which data centers account for the majority of disclosed revenue or assets are coded Undiversified Data Center, and all others are assigned by entity type. Their revenues are concentrated in data center leasing and related services, which makes them more dependent on data center demand. The remaining borrowers are classified as *Hyperscalers* (Amazon, Google, Meta, Microsoft, Apple, and Oracle), *Government*, or *Other*, which includes REITs, institutional investors, and other diversified corporations.[5] Table II, Panel A, reports summary statistics by borrower type.[6]

II. The Data Center Mortgage Market

The data show four features of this market. First, property-level data center mortgage lending is a recent phenomenon, expanding rapidly after mid-2021. Second, it consists of a small number of very large loans, many of them secured by several properties. Third, distribution-oriented lenders supply most of it. Fourth, most of the credit goes to Undiversified Data Center borrowers, while hyperscalers barely borrow through this channel.

Census Bureau data show that spending on data center construction is roughly flat from mid-2019 to mid-2021 and rises steadily afterward, while spending on other office buildings

[4]See Appendix Figure A2. *CoStar* does not observe GPU ownership, the operator's service contracts, or the ultimate compute customer.

[5]Hyperscalers are distinct from *CoStar*'s hyperscale facility type, which describes the building and is often owned by Undiversified Data Center borrowers.

[6]Appendix A provides complete variable definitions and data cleaning procedures. Appendix Table B.III lists the named borrowers in each category.

does not (Appendix Figure A3).[7] We date the boom to July 2021, when this rise begins. Mortgage lending follows the same pattern. Figure 1 documents cumulative gross data center mortgage originations from January 2015 by lender type and borrower type at origination. Through mid-2021 all lender types originate at modest rates. Origination then rises sharply.

Unlike residential mortgage markets, which generate millions of standardized originations, data center lending consists of a small number of large, bespoke transactions. After July 2021, 609 properties in our national sample receive mortgage financing. Many of these transactions are portfolio loans, in which a single loan is secured by several properties (Table I). Portfolio loans are more common at High-OTD lenders (Table II, Panel B).

Distribution-oriented (High-OTD) banks account for most of the growth in originations, while Low-OTD banks, non-bank lenders, and government lenders grow much more gradually (Figure 1, Panel A). High-OTD lenders finance all 49 loans of $1 billion or more and 92 of the 100 loans of $500 million or more. Most High-OTD mortgage credit is securitized, syndicated, or recorded under a trustee or agent, while most credit from other lenders is recorded under a single lender (Appendix Table B.IV). Panel B of Figure 1 shows that most of the lending boom goes to Undiversified Data Center borrowers. Since July 2021, the five largest borrowers account for almost half of mortgage credit. Undiversified Data Center borrowers carry approximately three times the median mortgage debt per square foot of all other borrowers (Table II, Panel A). Hyperscalers own 27% of data center space in our sample but account for 0.1% of mortgage debt.

Undiversified Data Center borrowers predominantly own colocation buildings, which sell capacity to many customers rather than lease the building to a single tenant.[8] Their revenues

[7] The timing also coincides with public announcements of large private investments in data centers, beginning with Blackstone's acquisition of QTS in June 2021 (Blackstone press release, June 7, 2021), and with a jump in data center securitization that year (CRE Finance Council, Data Center E-Primer, 2026).

[8] *CoStar* records the real-estate tenant or operator, not its customers. Hyperscalers lease the entire building at 6% of mortgaged properties, mostly owned by Other borrowers. They may also contract for part of the colocation capacity. While we do not observe this in our dataset, Digital Realty, for example, reports it has more than 5,000 customers, including Meta Platforms, and its largest accounts for about 12% of annualized recurring revenue (Form 10-K, 2025).

therefore rest on many customer contracts rather than on a single long-term lease, which leaves them exposed to changes in data center demand.

III. Empirical Strategy

Our empirical strategy is a difference-in-differences (DiD) design that compares High-OTD and other lenders before and after the expansion in data center construction that begins in July 2021, both in the leverage of the properties they finance and in the composition of their borrowers. The identifying assumption is that, absent the boom, the gap between High-OTD and other lenders would have remained stable. The event studies show no statistically significant differential pre-trends.

The event studies report periods 1 to 4 after the boom, periods -2 and -1 before it, and a single "All Prior" estimate pooling earlier observations.[9] The event-study specification is

$$Y_{it} = \beta_H \,\mathrm{HighOTD}_i + \sum_{p \neq -1} \beta_p \,(\mathrm{HighOTD}_i \times \mathbf{1}\{t \in p\}) + \mathbf{W}'_{it}\lambda + \gamma_s + \delta_t + \eta_p + \varepsilon_{it}, \qquad (1)$$

where p indexes these periods and period -1 is omitted in estimation. The figures plot each β_p relative to the average of the three pre-boom points, so the period -1 point equals minus that average. The vector $\mathbf{W}_{it}$ contains tenancy type and local macroeconomic controls, and γ_s, δ_t, and η_p denote state, loan-origination-year, and event-period fixed effects. We estimate the specification for ln(LoanAmount), ln(DebtPSF), ln(DebtToRent), and ln(LTV).

The corresponding static DiD specification is

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

[9] Event time is divided into 15-month periods anchored at July 2021. Earlier observations are pooled because individual periods in the long pre-boom tail contain few loans. The data extend through September 10, 2026. For deal-count outcomes, boundary-period counts are expressed per 15 months using their exact calendar exposure.

where Post_t equals one for loans originated in July 2021 or later. The vector $\mathbf{X}_{it}$ adds log rentable building area and local power supply and electricity prices to the event-study controls. Each property appears once, and β_{DiD} measures the differential post-boom change at High-OTD lenders. Standard errors are clustered by lender throughout.

The borrower-allocation analysis is estimated at the lender-period level:

$$N_{\ell p} = \beta_{\text{DiD}} \left(\text{HighOTD}_\ell \times \text{Post}_p\right) + \alpha_\ell + \eta_p + \varepsilon_{\ell p}, \tag{3}$$

where $N_{\ell p}$ is the number of new loans and refinancings in the properties' loan histories that lender ℓ originates for the indicated borrower group, expressed per 15 months, and α_ℓ and η_p are lender and period fixed effects. The panel includes every lender in every fifteen-month period, with zeros for periods in which a lender finances no deals. The event-study version is the same except that HighOTD is interacted with each period, and the earliest pre-boom periods share the "All Prior" interaction.

IV. Leverage and Borrower Composition at High-OTD Lenders

This section documents how mortgage origination by High-OTD lenders changes after July 2021 on two margins. On the intensive margin, we measure the size of the loans they originate, using the loan amount per property and debt relative to real estate (Section IV.A). On the extensive margin, we measure which borrowers receive these loans, using the number of deals with Undiversified Data Center and other borrowers (Section IV.B). Both margins point the same way, toward more debt against the same real estate and toward borrowers whose revenues depend more on data center demand.

A. Intensive Margin: Loan Size and Debt Relative to Real Estate

Figure 2 reports the event-study estimates for loan amounts per property (Panel A) and for debt per square foot, debt-to-rent, and loan-to-value (Panels B to D), which divide the loan amount per property by the size, rent, and value of the financed property. None shows statistically significant differential pre-trends (joint tests in Appendix Table C.I), and all show a large, persistent High-OTD gap after July 2021. Appendix Figure C1 shows the same patterns in levels.

Table III quantifies the comparison. The ratio of High-OTD to other-lender loan size rises close to fivefold (column 1), with no statistically detectable increase in the size of the properties they finance (column 5).[10] The ratio of High-OTD to other-lender debt relative to real estate rises between five- and sixfold across the three measures (columns 2 to 4).

B. Extensive Margin: Borrower Composition

Figure 3 shows no statistically significant differential pre-trends (joint test in Appendix Table C.I) and a post-boom divergence between High-OTD and other lenders in deals with Undiversified Data Center borrowers. Appendix Figure C2 shows the same pattern in levels. Table IV shows that the High-OTD gap in Undiversified Data Center deals widens after July 2021, much more than the gap in deals with Other borrowers, and the difference is statistically significant. Before July 2021, High-OTD lenders finance about 4 times as many Undiversified Data Center deals per lender as other lenders, and after July 2021 about 21 times as many. For Other borrowers, the ratio stays at about 4.

Table V estimates the property-level specification separately by borrower type. After July 2021, the increase in loan size and leverage at High-OTD lenders is larger for Undiversified Data Center borrowers than for other borrowers. These borrowers also take larger loans in general (Table II, Panel A), so the shift toward them adds to the larger loans. High-OTD

[10]The loan-size premium remains positive and significant for single-property loans, where the loan amount per property equals the total loan amount.

lenders therefore place more debt on the same real estate and direct most of it to borrowers whose revenues depend more on data center demand.

V. Alternative Explanations

This section examines specific alternative explanations for the results. We first examine explanations for the increase in leverage, based on property characteristics (Section V.A), technology (Section V.B), the composition of borrowers and lenders (Section V.C), and lender-wide changes (Section V.D). We then turn to checks that cover both margins, on lender size and classification (Section V.E) and on the cutoff date and the construction of the outcomes and sample (Section V.F).

A. Property Characteristics

A first concern is that High-OTD lenders financed different properties after 2021 and that the premium reflects this compositional shift. The characteristics of the financed properties do not change differentially. Appendix Table D.I, Panel A, estimates the High-OTD $\times$ Post-Jul'21 specification with each property characteristic as the outcome, including property size, construction status, assessed values per square foot, and building quality. None changes significantly.[11] Reweighting Low-OTD bank loans to resemble High-OTD loans in property size, location, origination year, tenancy, local power supply and electricity prices, and local economic conditions gives similar estimates (Appendix Table D.II).

B. Technology

We measure debt relative to observable real estate because the value of data center equipment is not observed at the property level and is highly uncertain over the medium term

[11] After 2021, High-OTD lenders finance a larger share of colocation facilities. Controlling for data center facility type (enterprise, hyperscale, or colocation) leaves the estimates similar, as does restricting the sample to colocation facilities, the largest group.

(Section I.C). A natural question is whether High-OTD lenders finance more technology-intensive facilities, with more equipment per square foot, so that the larger loans finance equipment rather than more debt against the same real estate. Panel B of Appendix Table D.I shows that power capacity, power density, tier, and redundancy do not change differentially, and controlling for them leaves the estimates similar (Appendix Table D.III). Debt per kW of utility power capacity also rises at High-OTD lenders.

C. Borrowers and Lenders

A related concern is that High-OTD lenders lent to different borrowers after 2021. Appendix Table D.IV compares loans within the same borrower and within the same lender. Borrower fixed effects compare each borrower's own loans before and after the boom, borrower-by-event-period fixed effects compare the same borrower's loans from different lender types within a 15-month period, and lender fixed effects compare each lender with itself. The premium in loan amount per property and in debt per square foot remains positive and statistically significant in all three. Restricting the sample to borrowers that already borrowed from any lender before July 2021, in *CoStar* or *DealScan*, yields larger High-OTD coefficients (column 4).

D. Lender-Wide Changes: Comparison with Warehouses

We next ask whether the premium reflects a change in High-OTD lenders' lending in general or something specific to data centers. Something other than data centers may have changed for High-OTD lenders after 2021, such as cheaper funding, more balance-sheet capacity, or a greater appetite for risk. A change common to all of their lending would also raise their loan amounts per property and debt measures on other real estate, not only on data centers. We test this with warehouses, the closest real-estate counterpart to data centers. Warehouses are large industrial buildings financed by many of the same lenders, with 130 lenders, 28 of them High-OTD, lending in both sectors.

We stack data center and warehouse loans originated since January 2015 and estimate

$$Y_{it} = \beta_3 \left(\text{HighOTD}_i \times \text{Post}_t \times \text{DC}_i\right) + \beta_2 \left(\text{HighOTD}_i \times \text{Post}_t\right) + \beta_1 \left(\text{HighOTD}_i \times \text{DC}_i\right) + \beta_{PD} \left(\text{Post}_t \times \text{DC}_i\right) + \beta_H \, \text{HighOTD}_i + \beta_P \, \text{Post}_t + \mathbf{X}'_{it}\theta + \gamma_{s,k} + \delta_{t,k} + \varepsilon_{it}, \tag{4}$$

where DC_i equals one for data centers and zero for warehouses, and $\gamma_{s,k}$ and $\delta_{t,k}$ are sector-by-state and sector-by-origination-year fixed effects. The coefficient β_3 measures the post-boom change at High-OTD lenders on data centers relative to warehouses, and β_2 the change on warehouses.

Appendix Table D.V reports the estimates. Comparing sectors removes changes common to a lender's whole book, and comparing lenders removes the general rise in data center financing. High-OTD lenders do not make larger loans or raise their debt measures on warehouses after July 2021, and their data center premium relative to warehouses is close to the baseline for loan amount per property and all three debt measures (Panel A). The central comparison is Panel B, which adds lender-by-year fixed effects. It compares the same lender's data center and warehouse loans in the same year and absorbs conditions common to that lender's real-estate lending in that year, such as its funding conditions and appetite for risk. The premium holds in this comparison.[12] The comparison also addresses changes over time in how *CoStar* records loans, since such changes would also affect warehouse loans by the same lenders.

E. Lender Size and OTD Classification

The results could reflect bank size rather than distribution, and they could depend on how we classify lenders. Controlling for predetermined bank size leaves the estimates similar, and bank size itself does not predict the post-boom change (Appendix Table D.VI, Panel

[12]Restricting the sample to lenders that lend in both sectors gives similar estimates, and restricting it to warehouses at least as large as the smaller data centers gives larger estimates. Year-by-year estimates show no significant differential pre-trends between the two sectors.

A). Excluding the 5, 10, or 15 largest banks by total assets leaves the estimates positive and significant (Panel B). The estimates also remain positive and significant under alternative classifications: a 2022 to 2025 window, excluding trustee-recorded loans or foreign banks, and a continuous OTD score (Appendix Table D.VII).

F. Cutoff Date, Outcomes, and Sample

We date the boom to July 2021, when data center construction spending turns up (Appendix Figure A3). Shifting the cutoff three or six months earlier or later leaves all estimates positive, similar in magnitude, and significant (Appendix Table D.VIII), as does restricting the sample to loans originated since 2015. Restricting the sample to single-property loans, which need no allocation across properties, gives smaller but significant estimates (Appendix Table D.IX, Panel A). Excluding non-operational properties or refinancings, or controlling for refinancing, leaves the estimates positive and significant (Panels B to D). The borrower-allocation estimate is similar with alternative count measures (Appendix Table D.X), and winsorizing, trimming, and alternative clustering leave all estimates significant.

VI. Interpretation and Implications

A. Interpretation

Our comparisons narrow the possible explanations. The larger loans are unlikely to reflect larger or different properties, since property characteristics do not change differentially and reweighting on them gives similar estimates. They are unlikely to reflect more equipment-intensive facilities, since power capacity, density, and tier do not change differentially. Differences in borrower composition are also unlikely to explain them, since the same borrower in the same period receives larger loans from High-OTD lenders, and the estimates are larger for borrowers that already borrowed before the boom. The premium also holds when compared with the same lender's warehouse loans in the same year, including large

warehouses, a comparison that absorbs conditions common to the lender in that year. Nor does bank size explain them.

The results show two changes in how High-OTD lenders lend against data centers after 2021. High-OTD lenders finance many more deals with Undiversified Data Center borrowers (Table IV), and the ratio of their debt relative to real estate to that of other lenders rises five- to sixfold (Table III). Three explanations could account for these changes.

Weaker underwriting incentives could produce both patterns. Lenders that distribute their loans have weaker incentives to screen and monitor borrowers (Gorton and Pennacchi, 1995; Keys et al., 2010). Weaker screening could allow more debt against the same real estate and more credit to borrowers that depend on a single source of demand.

Greater financing capacity and investor demand could produce both the larger loans and the shift toward sector-dependent borrowers. Distribution frees balance-sheet space (Loutskina, 2011). Controls for bank size (Appendix Table D.VI) and the comparison with the same lender's warehouse loans in the same year (Appendix Table D.V) address capacity common to a lender's whole book. Investor demand for data center debt could lower financing costs (Gete and Reher, 2021) and direct more credit toward sector-dependent operators.

Specialization and selection could also produce both patterns. Lenders specialize in industries where they have expertise (Blickle, Parlatore, and Saunders, 2026). High-OTD lenders already financed more deals with Undiversified Data Center borrowers before the boom (Appendix Figure C2), without a difference in leverage (Figure 2). Their expertise may have become more valuable as the market expanded and transactions became larger and more complex. Even within the same borrower, lenders may finance projects with different tenant contracts, guarantees, or development prospects.

B. Implications

These patterns do not depend on which explanation holds. After 2021, more debt rests on the same real estate, and it is owed mainly by borrowers whose revenues depend more on

data center demand. If that demand weakens and other repayment sources fall short, the real estate would cover a smaller share of the debt. Loan performance would help separate these explanations, but data center mortgages have not yet been through a downturn. We also do not observe retained exposures, so we cannot establish where the risk ultimately resides.

VII. Conclusions

Our new property-level database shows that the mortgage financing of the AI data center boom is concentrated among a few distribution-oriented lenders. As the market expanded, these lenders extended more debt against comparable real estate and shifted credit toward operators whose revenues depend on data center demand.

Part of the AI buildout therefore rests on mortgage debt that is large relative to the underlying real estate and owed by sector-dependent borrowers. How these loans perform through the market's first downturn, and where the associated credit losses ultimately fall, remain important questions.

References

Adrian, Tobias and Hyun Song Shin (2010). Liquidity and leverage. *Journal of Financial Intermediation 19*(3), 418–437.

Benmelech, Efraim, Jennifer Dlugosz, and Victoria Ivashina (2012). Securitization without adverse selection: The case of CLOs. *Journal of Financial Economics 106*(1), 91–113.

Black, Lamont K., Chenghuan Sean Chu, Andrew Cohen, and Joseph B. Nichols (2012). Differences across originators in CMBS loan underwriting. *Journal of Financial Services Research 42*(1–2), 115–134.

Blickle, Kristian, Cecilia Parlatore, and Anthony Saunders (2026). Specialization in banking. *Journal of Finance 81*(3), 1531–1572.

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

Buchak, Greg, Gregor Matvos, Tomasz Piskorski, and Amit Seru (2024). Beyond the balance sheet model of banking: Implications for bank regulation and monetary policy. *Journal of Political Economy 132*(2), 616–693.

Darmouni, Olivier, Clemens Lehner, and Yuqi Zhang (2026). The energy cost of AI and data centers. Working Paper, SSRN 6751399.

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

Gargano, Antonio and Marco Giacoletti (2025, November). Subsidizing the cloud: U.S. state incentives to data centers. Working paper.

Gete, Pedro and Michael Reher (2021). Mortgage securitization and shadow bank lending. *Review of Financial Studies 34*(5), 2236–2274.

Ghent, Andra C. and Rossen Valkanov (2016). Comparing securitized and balance sheet loans: Size matters. *Management Science 62*(10), 2784–2803.

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

Heiss, Alexander, Zacharias Sautner, and Thomas Schmid (2026, August). Powering AI: How do data centers affect renewable energy investment? Working paper.

Irani, Rustom M., Rajkamal Iyer, Ralf R. Meisenzahl, and José-Luis Peydró (2021). The rise of shadow banking: Evidence from capital regulation. *Review of Financial Studies 34*(5), 2181–2235.

Keys, Benjamin J., Tanmoy Mukherjee, Amit Seru, and Vikrant Vig (2010). Did securitization lead to lax screening? evidence from subprime loans. *Quarterly Journal of Economics 125*(1), 307–362.

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

Mian, Atif and Amir Sufi (2009). The consequences of mortgage credit expansion: Evidence from the U.S. mortgage default crisis. *Quarterly Journal of Economics 124*(4), 1449–1496.

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

Rand, Joseph, Anna Cheyette, Chris Talley, Steven Zhang, Will Gorman, Ryan Wiser, Joachim Seel, Seongeun Jeong, and Fritz Kahrl (2026, June). Queued up: 2026 edition, characteristics of power plants seeking transmission interconnection as of the end of 2025. Technical report, Lawrence Berkeley National Laboratory.

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

Van Nieuwerburgh, Stijn (2026). Financing the AI buildout. NBER Working Paper No. 35865, forthcoming in Brookings Papers on Economic Activity, Fall 2026.

Wachter, Jessica A. and Jonathan Wachter (2026). What investment data implies about the ai transition. NBER Working Paper 35290, National Bureau of Economic Research.

Yue, Daniel and Yiyang Zeng (2026, March). The local economic effects of data center entry. Working paper, SSRN 6497238.

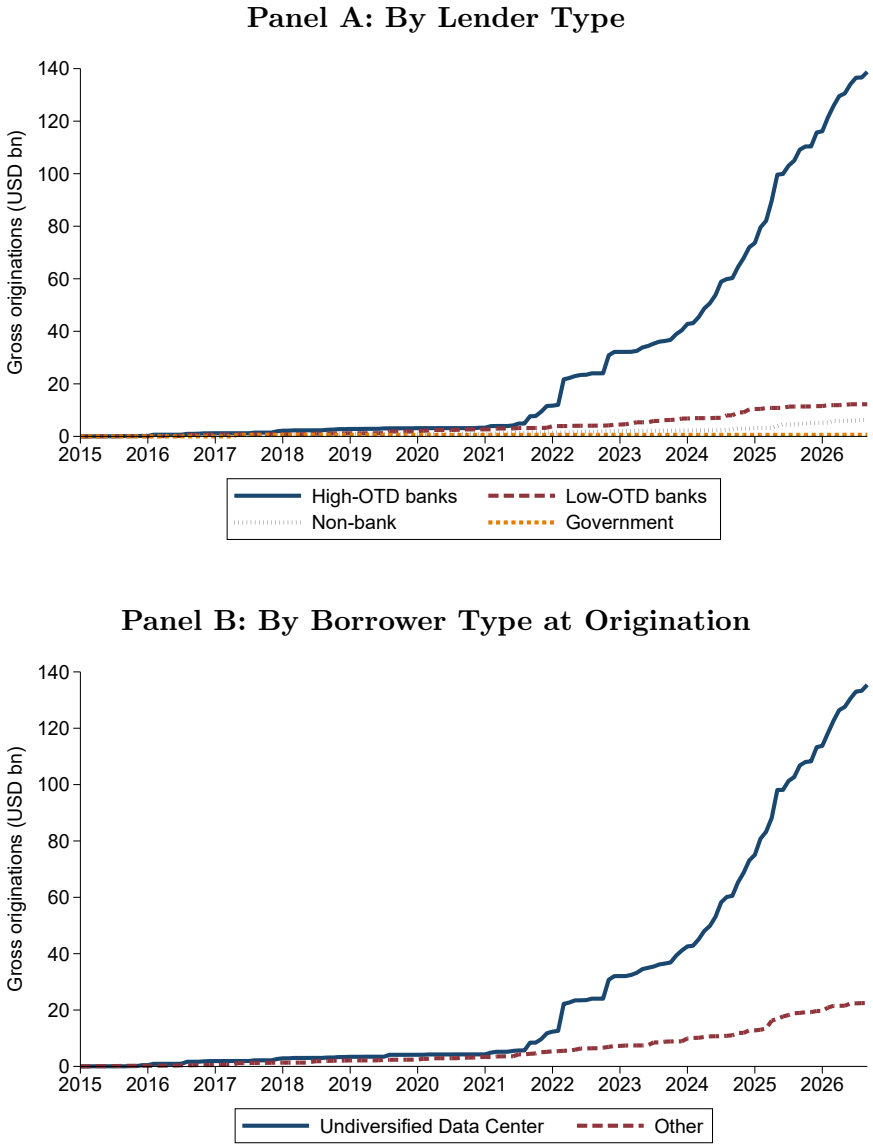


Figure 1. Cumulative Gross Mortgage Originations

Total mortgage debt originated since January 2015, cumulated by month and including refinancings. Panel A shows it by lender type and Panel B by borrower type at origination.

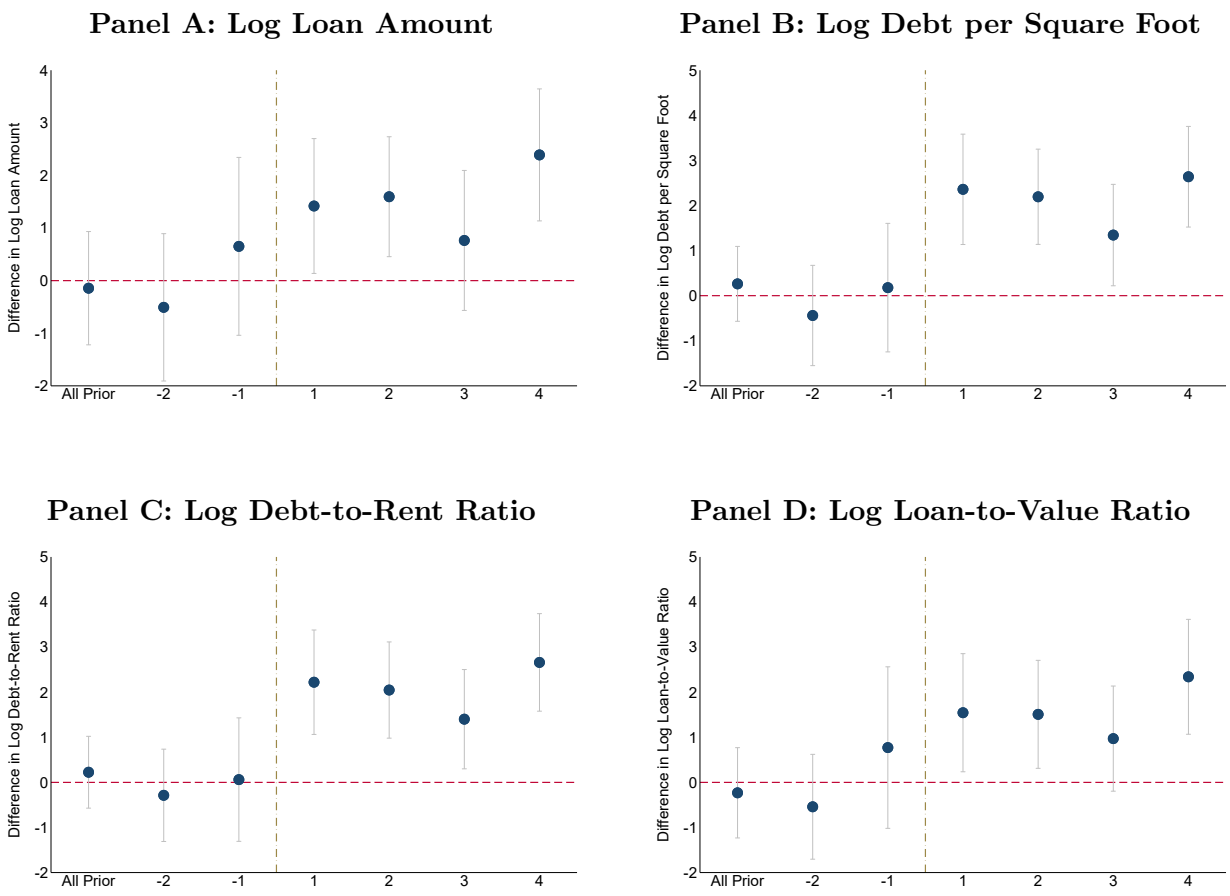


Figure 2. Loan Size and Debt Relative to Real Estate

Event-study estimates of Equation 1 for the loan amount per property and each debt measure. Standard errors are clustered by lender, and error bars are 95% confidence intervals.

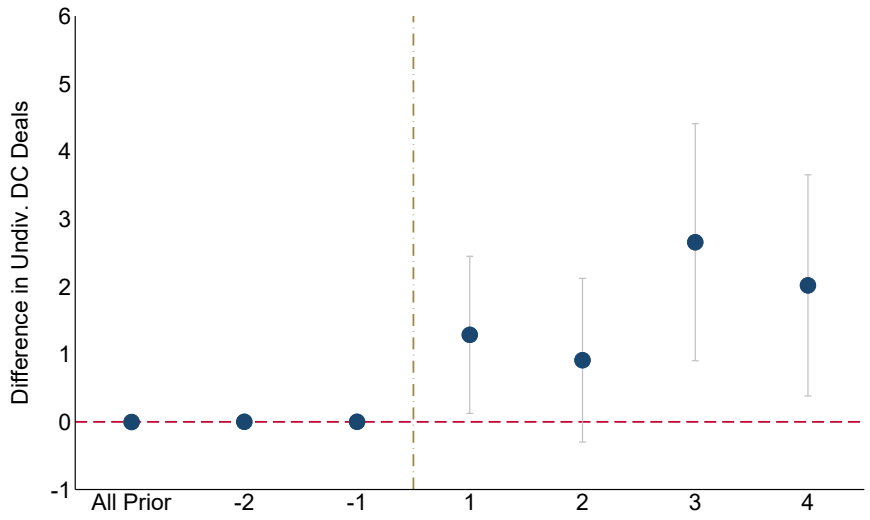


Figure 3. Deals with Undiversified Data Center Borrowers

Estimates of the event-study version of Equation 3. Standard errors are clustered by lender, and error bars are 95% confidence intervals.

Table I. Summary Statistics, 1992–2026

Variable	N	Median	SD	Min	Max
Panel A: Property characteristics					
Rentable building area (SF)	854	150,860	379,432	1,125	8,600,000
Rent per square foot (USD)	854	13.6	5.7	1.0	50.7
Panel B: Loans and debt					
Loan amount (USD M)	854	26.3	294.0	0.0	3,500.0
Properties securing the loan	854	1.0	9.1	1.0	147.0
Debt per square foot (USD)	854	223.2	2,731.5	0.1	36,896.6
Debt-to-rent ratio	854	16.4	234.0	0.0	3,113.9
Loan-to-value ratio	854	0.3	6.5	0.0	59.4
Years since origination	854	3.0	6.6	0.0	34.0

Note: RBA denotes rentable building area. Loan amount is the amount allocated to each property, using documented property allocations where available and RBA or documented collateral shares otherwise. Properties securing the loan is the documented number of collateral properties of each property's loan, or the number of sample properties that share the loan when it is not documented. Years since origination is measured as of 2026.

Table II. Borrowers, Lenders, and the OTD Classification

Panel A: Borrower type

	Undiversified Data Center	Other	Difference
Properties	506	348	
Median loan amount (USD M)	87.2	9.9	77.3
Median property RBA (SF)	205,118.0	101,956.0	103,162.0
Median debt per square foot (USD)	378.3	116.8	261.5

Panel B: Lender type

	High-OTD	Other	Difference
Properties	452	402	
Median loan amount (USD M)	101.2	7.2	93.9
Median property RBA (SF)	190,000.0	118,498.0	71,502.0
Median debt per square foot (USD)	560.4	79.5	481.0
Properties in multi-property loans (share)	0.60	0.33	0.27

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

	High-OTD	Low-OTD	Difference
Banks	13	100	
Loan assets/assets	52.91	72.22	−19.31
Repo/assets	9.44	0.37	9.07
Trading assets/assets	4.68	0.09	4.58

Note: RBA denotes rentable building area. Panels A and B report property-level counts, medians, and shares. In Panel B, Other lenders are Low-OTD banks, non-bank lenders, and government lenders, and multi-property loans are loans secured by more than one property. Panel C reports bank-level means of the regulatory ratios underlying the OTD score, restricted to banks classified using their own *Call Reports*.

Table III. Loan Size and Debt Relative to Real Estate

	ln(LoanAmount)	ln(DebtPSF)	ln(DebtToRent)	ln(LTV)	ln(RBA)
	(1)	(2)	(3)	(4)	(5)
High-OTD × Post-Jul'21	1.613***	1.769***	1.770***	1.655***	−0.284
	(0.395)	(0.354)	(0.342)	(0.373)	(0.213)
Controls and FE	Yes	Yes	Yes	Yes	Yes
R^2	0.514	0.494	0.482	0.474	0.322
Observations	854	854	854	854	854

Note: This table reports estimates of Equation 2 with controls for local macroeconomic conditions, tenancy type, and local power supply and electricity prices, and state and origination-year fixed effects. Columns 2 to 4 also control for RBA, the rentable building area that measures property size. In column 1, the outcome is the loan amount per property, using documented property allocations where available and RBA or documented collateral shares otherwise. In column 5, the outcome is RBA. Standard errors clustered at the lender level are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table IV. OTD and Borrower Allocation

	Undiv. DC Deals	Other Deals	Difference
	(1)	(2)	(3)
High-OTD × Post-Jul'21	1.720***	0.283**	1.437***
	(0.610)	(0.134)	(0.545)
Lender and period FE	Yes	Yes	Yes
R^2	0.190	0.113	0.176
Lender-periods	10,395	10,395	20,790

Note: This table reports estimates of Equation 3. The outcome is the number of new loans and refinancings in the properties' loan histories made to Undiversified Data Center borrowers (column 1) or Other borrowers (column 2). Column 3 is the gap between columns 1 and 2, estimated in a single regression on both counts. Standard errors clustered at the lender level are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table V. Loan Amount and Debt by Borrower Type

	ln(LoanAmount) (1)	ln(DebtPSF) (2)	ln(DebtToRent) (3)	ln(LTV) (4)
Panel A: Undiversified Data Center borrowers				
High-OTD × Post-Jul'21	2.517***	2.558***	2.446***	2.854***
	(0.818)	(0.776)	(0.763)	(0.757)
R^2	0.482	0.478	0.479	0.466
Observations	506	506	506	506
Panel B: Other borrowers				
High-OTD × Post-Jul'21	0.673	1.012**	0.969**	0.803
	(0.470)	(0.462)	(0.461)	(0.534)
R^2	0.519	0.530	0.499	0.498
Observations	348	348	348	348
Difference (Panel A − Panel B)				
High-OTD × Post-Jul'21	1.844**	1.546*	1.477	2.051**
	(0.935)	(0.910)	(0.901)	(0.911)
R^2	0.551	0.523	0.512	0.516
Observations	854	854	854	854
Controls and FE	Yes	Yes	Yes	Yes

Note: This table reports estimates of Equation 2 separately for Undiversified Data Center borrowers (Panel A) and other borrowers (Panel B), with the controls of Table III and event-period instead of origination-year fixed effects. Column 1 omits RBA from the controls. The Difference row is the gap between Panels A and B, estimated in a single regression on both groups of borrowers. 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

A.1 Data Sources and Sample Construction

***CoStar* data center inventory.** The starting point is the *CoStar* national inventory of U.S. data center properties as of September 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 I reports summary statistics for the mortgaged sample.

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. Records with unverified amounts or financing unrelated to the data center are excluded. Loans secured by several properties are allocated across them as described in Appendix A.5.

Lender OTD measures. We link lenders to regulatory *Call Report* data to obtain OTD-based distribution orientation measures. Appendix A.3 describes the OTD score.

***DealScan* syndicated loan data.** We match *DealScan* lender-tranche records to our own lenders and examine their syndication behavior across all of their *DealScan* facilities, not only data center–related ones, to compare the lending behavior of High- and Low-OTD banks (Appendix A.3).

Local macro controls. County-level annual GDP, population, and unemployment rate come from the Bureau of Economic Analysis and Bureau of Labor Statistics and are 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. Local power supply and electricity prices use each property's own county.

Construction spending. Monthly value of private construction put in place for data centers and other office buildings, seasonally adjusted, comes from the Census Bureau's construction spending survey. It is used in Section II and Appendix Figure A3.

Data vintages. The *CoStar* data center inventory was accessed in September 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 December 2025 and cover 1959 through 2025 Q2. *DealScan* data were obtained in July 2026 and cover facilities through March 2026. *CoStar* warehouse data were obtained in October 2026. Census Bureau construction spending data were obtained in October 2026 and cover January 2014 through August 2026.

A.2 Estimation Samples

The baseline analysis uses all 854 mortgaged data center properties. Borrower heterogeneity analyses split this sample by borrower type. The borrower-allocation analysis instead counts every new loan and refinancing in the properties' loan histories, aggregated to lender-by-period observations. Some robustness exercises restrict it further, for example to single-property loans, operational properties, or bank lenders with a *Call Report* score. The warehouse comparison (Appendix Table D.V) uses a separate *CoStar* extract of warehouse properties in the states where the sample's data centers are located, with loans originated since January 2015.

A.3 Borrower and Lender Classification

Borrower types. We classify borrowers into four groups: Undiversified Data Center, Hyperscalers, Government, and Other. The classification uses standardized borrower names and borrower group labels derived from *CoStar* ownership fields. It is made at the level of the ultimate owner, and recourse and sponsor guarantees are not observed. Iron Mountain and Mapletree Industrial Trust, whose disclosed businesses extend well beyond data centers, are classified as Other.

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, and High- or Low-OTD status follows from the decile rule described below. Nonbank and government/cooperative lenders are assigned separately by lender type. Foreign banks, which lack usable *Call Report* data, are classified as described below.

Originate-to-distribute (OTD) measures. Using *Call Report* data for 2017–2020, we standardize each bank's loan-asset, trading-asset, and repurchase-agreement-asset ratios (each measured as a share of total assets) within year and combine them into the OTD score:

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

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 data center mortgage debt despite representing a small share of data center–lending banks.

Foreign banks. For foreign banks we lack usable parent-level *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 B.II), and is applied to Mizuho, Royal Bank of Canada, Sumitomo Mitsui, Société Générale, Natixis, CIBC, MUFG, Bank of Montreal, National Bank of Canada, Royal Bank of Scotland, BNP Paribas, ING, UBS, Toronto Dominion, Barclays, and Deutsche Bank. For Toronto Dominion, Barclays, and Deutsche Bank, the available *Call Reports* belong to U.S. subsidiaries rather than the parents. Eurohypo, absent from *DealScan*, is classified Low-OTD. 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 B.II, Panel A, are constructed

directly from *DealScan*'s own fields. Full syndication and club deal are indicators for *DealScan*'s distribution method. Lead arranger indicates a lead role (arranger, bookrunner, lead or mandated lead arranger, administrative or syndication agent).

Syndicated and securitized loans. Syndicated loans are classified by their first-listed bank. Securitized loans recorded under a trustee or servicer, and loans recorded under an agent for a lender group, are coded High-OTD. Non-bank lenders belong to the comparison group of other lenders, so no result we report requires imputing an OTD score.

Financing structure. We assign facilities to mutually exclusive financing structures in the following order. CMBS/ABS facilities appear in *CoStar*'s CMBS records or name a securitization. Agent- or trustee-recorded facilities are recorded under a securitization trustee or a loan agent for a lender group. Syndicated facilities list two or more lenders in one *CoStar* record or match a *DealScan* facility with two or more lenders, matched by borrower, lender, date, and amount and reviewed by hand. The remaining facilities have one recorded lender and no observed distribution.

Bank size variables. We also construct predetermined bank-size measures, using *Call Report* data for domestic banks and publicly available balance sheet data for foreign institutions, as robustness controls in Appendix Table D.VI: the log of total assets as of 2020 Q4 and its interaction with Post-Jul'21 (Panel A), and exclusions of the 5, 10, and 15 largest lenders ranked by 2020 consolidated assets (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 identity. 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; Harris Trust & Savings and Bank of Montreal; 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. We convert *CoStar*'s reported shell rent into a numeric rent per square foot measure, imputing missing or invalid values with the median rent per square foot among data centers in the same county, falling back to the state median. Rents and property characteristics are as of the September 2026 *CoStar* extract, not the origination date. For summary statistics and regressions, rent is interpreted as shell (base) rent and excludes colocation usage charges and other operating pass-throughs.

Local power supply and electricity prices. Local power supply is the generation capacity requested in the property's county interconnection queue over the three years up to the origination year, in total and for renewable and gas generation, from the Berkeley Lab interconnection queue data (Rand et al., 2026). Electricity prices are state average commercial and industrial retail prices from the Energy Information Administration. For 2026, the queue data end in 2025.

Building age. *CoStar* provides building age, the number of years since construction, mea-

sured as of 2026. 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 D.IX, Panel B, 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 Loan Amounts and Debt Measures

Debt allocation for multi-property loans. Loan origination records may correspond to a loan secured by multiple properties. To construct property-level loan amounts and debt measures, we use documented property allocations where available and otherwise apportion the balance across collateral properties by rentable building area (RBA), with equal shares when size is unavailable. 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

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

For each financing, property-level debt is its allocated balance. RBA-proportional allocation is the economically appropriate way to assign a portfolio loan across its pledged facilities. The property-level regressions use one loan per property, the most recent qualifying loan, with documented CMBS records taking precedence. The deal counts use every new loan and refinancing in the properties' loan histories. Counting a loan requires only its date, lender, and borrower, which the histories record for every loan. The leverage measures also require the size, rent, and value of the financed property, which *CoStar* reports as of the 2026

extract, so they use the most recent loan, whose property best matches the one observed.

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 is allocated property debt divided by annual shell rent, which is shell rent per square foot times RBA.

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 value using the county or state median assessed value per acre. 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 approximates the highest-value-use principle, since the land-value approach captures the value of the site in alternative use, and yields the most conservative loan-to-value ratio. 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. The warehouse sample of Section V.D lacks matched parcel records, so the comparison in Appendix Table D.V uses

$\max\{\text{DCF}_i, \text{Land}_i\}$ in both sectors.

Origination timing. We parse *CoStar*'s loan origination date and extract the origination year. Years since origination is defined as $2026 - \text{OriginationYear}_i$.

Refinancing. We use *CoStar* loan histories to distinguish refinancing of existing data center debt from new financing. Loans with no qualifying earlier data center debt, loans following repayment of earlier debt, and loans preceded only by debt associated with an earlier property use are classified as new originations. For completed buildings, earlier-use debt includes loans originated more than two years before data center construction and identified land or farm loans. Construction financing for proposed, planned, or under-construction data centers is retained.

Cumulative originations. Figure 1 and the deal counts use each property's full *CoStar* loan history, including refinancings, while the property-level regressions use only the most recent loan per property. A loan recorded on several properties, or recorded more than once, is counted once and allocated across its collateral properties as described above. When only the number of properties in a collateral pool is documented, each property receives an equal share. Borrower type at origination follows the borrower named in the loan record, then recorded sales, then the current owner. Loans recorded under an agent or title company with no documented creditor are grouped with Low-OTD banks, which if anything understates the High-OTD share.

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 reported in Table I (full collateral-value construction in Appendix Table A.II). This observation and a small number of similar cases reflect genuine single-property loans to specialized data center operators in which the real-estate collateral value, computed as the maximum of assess-

ment, 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. We found no evidence that they reflect data or allocation errors, and winsorizing or trimming extreme values leaves the main results unchanged.

Excluded records. *CoStar* occasionally attaches records to a data center property that do not finance the data center, such as home or farm loans to individuals on parcels later assembled into the site, loans to other borrowers sharing a collateral list, a loan between affiliates of the same owner, or a placeholder amount. These records are excluded. One Flexential record is also excluded. *CoStar* reports a $2.0 billion loan from an unidentified private lender on a 33,588 square foot facility, which matches no public financing, and we cannot verify it.

A.6 Sample Coverage and Representativeness

Appendix Table A.I compares properties with and without an observed mortgage record. A property without an observed mortgage record may be unlevered, financed through a channel *CoStar* does not capture, or not yet reflected in *CoStar*'s records, and these cases cannot be distinguished. The starting universe reflects *CoStar*'s September 2026 inventory, so properties removed from its active inventory before that date are excluded.

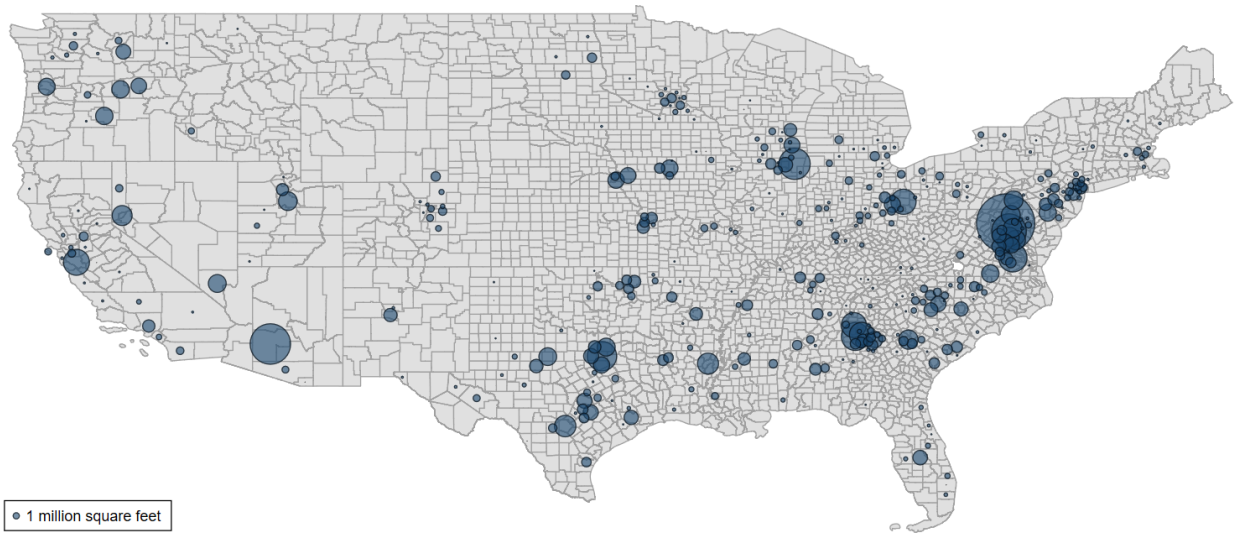


Figure A1. Data Center Capacity by County

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

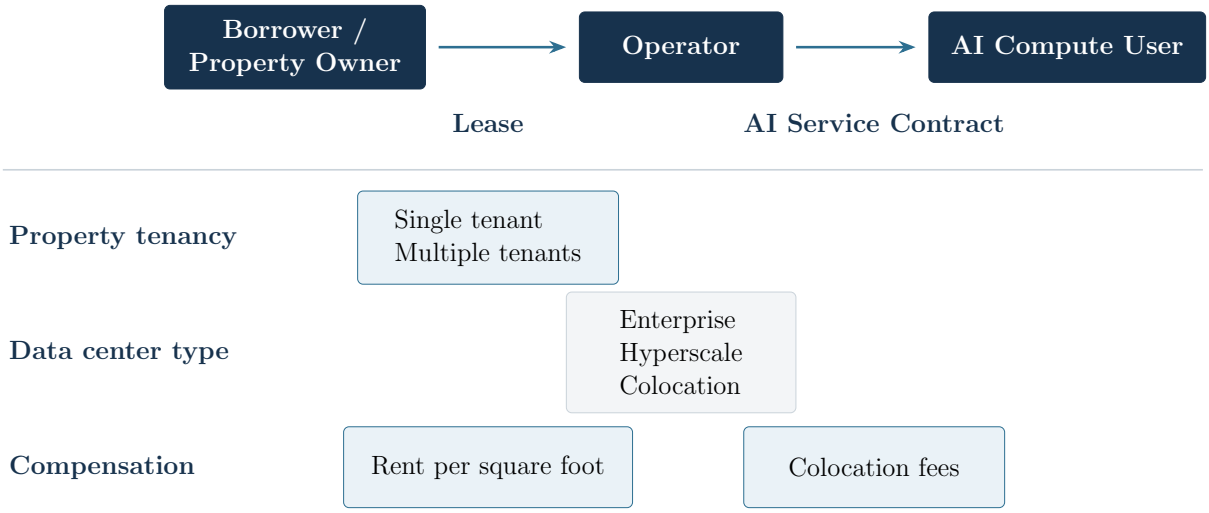


Figure A2. *CoStar* Variables Along the Data Center Supply Chain

CoStar reports the property owner, property tenancy (one tenant or several), rent per square foot, and a data center type category (Enterprise, Hyperscale, or Colocation). Also shown are the operator's AI service contracts with its customers and the associated colocation fees, which are not available in *CoStar*.

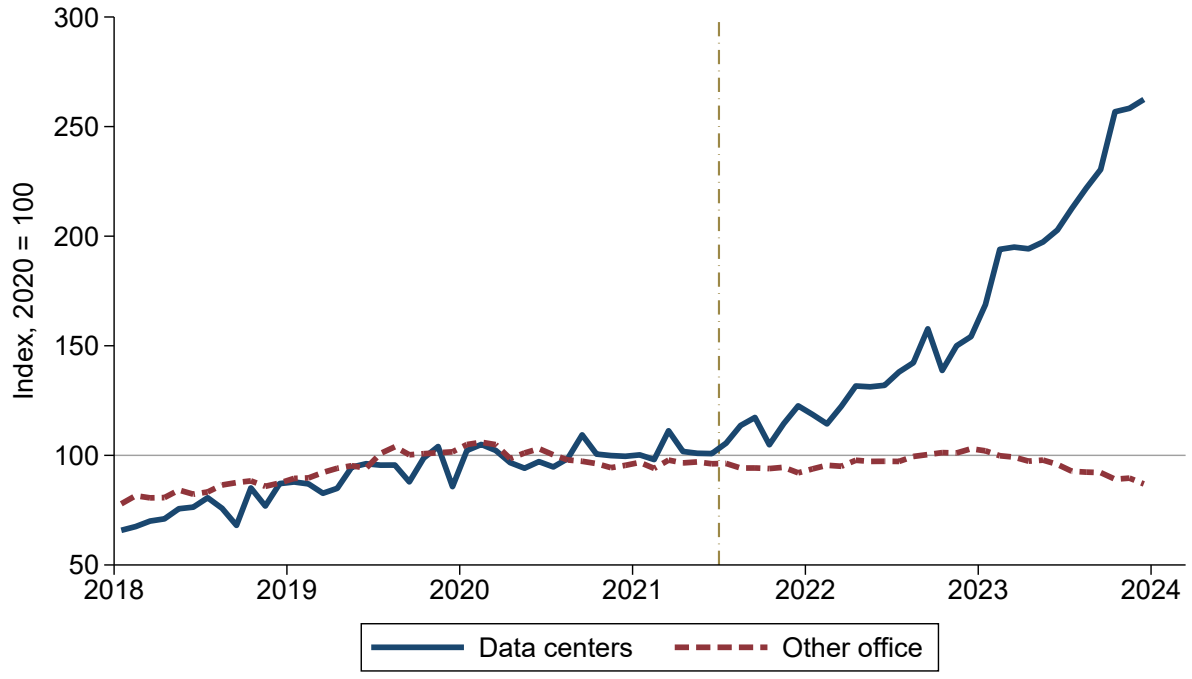


Figure A3. Construction Spending on Data Centers

Census Bureau value of private construction put in place, seasonally adjusted annual rate, for data centers and other office buildings. Each series is indexed to its 2020 average. The figure covers January 2018 to December 2023, and the vertical line marks July 2021.

Table A.I. Sample Coverage: Properties With vs. Without Observed Mortgage Debt

	No debt	Debt	Difference	Observations
RBA (sq ft)	295,273	232,801	−62,472***	2,933
Acreage	100.5	30.1	−70.4***	2,764
Building age	15.6	23.4	7.8***	1,931
Rent per sq ft (USD)	12.8	14.8	2.1***	2,166
Undiversified Data Center (share)	0.344	0.593	0.248***	2,933
Hyperscalers (share)	0.288	0.023	−0.264***	2,933
Existing building (share)	0.552	0.738	0.186***	2,933

Note: The table compares properties with a usable data center mortgage record (Debt) with those without one (No debt). Difference is the coefficient from a regression of each characteristic on an observed-debt indicator, with robust standard errors. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A.II. Collateral Values and Implied Loan-to-Value Ratios

Variable	N	Mean	SD	Min	Max
Panel A: Collateral and debt per square foot (USD)					
Tax assessment value per square foot	676	464.9	1,435.2	10.0	29,574.2
Discounted cash flow value per square foot	719	519.9	209.0	35.1	1,778.1
Acre-based land value per square foot	854	662.8	2,651.5	2.5	65,116.9
Real-estate collateral value per square foot	854	990.9	2,838.8	27.9	65,116.9
Debt per square foot	854	996.5	2,731.5	0.1	36,896.6
Panel B: Implied loan-to-value ratios					
Tax assessment basis	676	17.1	59.0	0.0	677.0
Discounted cash flow basis	719	2.6	7.2	0.0	88.8
Acre-based land value basis	854	8.2	25.6	0.0	422.2
Real-estate collateral value	854	2.4	6.5	0.0	59.4

Note: The tax assessment rows use only properties with a matched parcel assessment. Real-estate collateral value is the highest of the three values (Appendix A.5).

Appendix B: Borrower and Lender Classification

Table B.I. Lender Classification

High-OTD lenders	Low-OTD lenders	Non-bank lenders	Government lenders
Wilmington Trust (CMBS trustee)	Citizens Bank	Marquee Funding Group	Compeer Financial
Toronto Dominion Bank	U.S. Bank	Silver Point Finance	State of Ohio
Wells Fargo	Truist Bank	Individual lender	Badgerland Financial FLCA
Sumitomo Mitsui Banking	First Citizens Bank	Acore Capital Mortgage	Small Business Administration
Citibank	First Bank	American General Life Insurance	Farm Credit Mid-America FLCA
Goldman Sachs	Sunflower Bank	Athene Annuity & Life	AgGeorgia Farm Credit ACA
Bank of America	KeyBank	Split Oak Venture	County of Washington
JPMorgan Chase	Huntington National Bank	Roserock Capital Fund I	CoBank ACB
Computershare (CMBS trustee)	Capital One	Principal Life Insurance	
GLAS (loan agent)	Suntrust Bank	Argentic Real Estate Finance	
Société Générale	Comerica Bank	Column Financial	
Alter Domus US (loan agent)	Arizona Business Bank	Fidelity Guaranty Life Insurance	
Fifth Third Bank	Interbank	Griffin Capital Real Estate	
Morgan Stanley	UBS	Various other non-bank lenders	
Barclays	Pinnacle Bank		
Deutsche Bank	Bankwell Bank		
UMB Bank	Eagle Bank		
Royal Bank of Canada	Valley National Bank		
CIT Group	Ameris Bank		
CIBC	Various other Low-OTD lenders		
Bank of Montreal			
Natixis New York Branch			
Mizuho Bank			
Compass Bank			
MUFG Bank			
Merrill Lynch			
Uinta Bank			
Watermark Bank			
National Bank of Canada			
ING			
Cortland Capital Market Services (loan agent)			
BOK Financial			
Horizon Bank			
Lehman Brothers Holdings			
BNP Paribas			
Commerce Bank			
Royal Bank of Scotland			

Note: Lenders within each classification are sorted by number of matched properties in the sample, descending. Securitized loans recorded under a CMBS trustee, and loans recorded under a loan agent, are coded High-OTD and listed under the trustee's or agent's name (Appendix A.3).

Table B.II. *DealScan* Comparison of Lender Behavior

Panel A: Deal-level syndication behavior			
	Full Syndication	Club Deal	Lead Arranger
High-OTD bank	−0.035***	0.031***	0.224**
	(0.009)	(0.009)	(0.083)
Log deal size	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
R^2	0.034	0.035	0.051
Observations	4,798	4,798	4,798
Panel B: Bank-level scale			
	High-OTD	Low-OTD	Difference
Mean *DealScan* volume per bank ($bn)	430.6	34.4	396.2***

Note: The sample is the *DealScan* lender positions of the banks in our lender sample that have an OTD classification. *DealScan* covers syndicated facilities and is not restricted to mortgages. Panel A regresses each behavior indicator on a High-OTD indicator, log deal size, and year fixed effects, with Low-OTD banks as the comparison group. Panel B reports mean *DealScan* deal volume per bank, and the difference uses an unequal-variance test. Standard errors clustered at the lender level are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table B.III. Borrower Classification

Undiversified Data Center		**Hyperscalers**	**Other**
STACK Infrastructure	5C Data Centers	Amazon (AWS)	Landmark Dividend
QTS	Ark Data Centers	Microsoft	Blackstone REIT
Switch, Inc.	BlackChamber Group	Google	Mapletree Industrial Trust
CyrusOne	Centra Digital	Meta	GI Partners
Digital Realty Trust, Inc.	DigiCo Infrastructure REIT		Menlo Equities
EdgeConneX	Red Wolf DCD Properties LLC		Iron Mountain Incorporated
Vantage Data Centers	ScaleMatrix		CleanSpark, Inc.
Aligned Data Centers	Sentinel Data Centers, LLC		Igneo Infrastructure Partners
CSquare	TRG Datacenters		Core Scientific, Inc.
DataBank	AiNET		BGRE
Compass Datacenters	Apotech Data Centers		Oppidan
DC BLOX	CDS Global Cloud		Penn Hurst Holdings DE LLC
Cologix	Cirrascale Cloud Services		Starwood Capital Group
Flexential	Coloswitch LLC		Aphorio Carter
Prime Data Centers	DP Facilities, Inc.		Blue Owl Capital
H5 Data Centers	DartPoints		StratCap
Serverfarm	Data Center Leasing, Inc.		Waha Inc.
Stream U.S. Data Centers, L.L.C.	Datacate, Inc.		Bouwinvest Real Estate Investors B.V.
TierPoint	Direct LTx		Innovation 411 Fee Owner LLC
CoreSite	Eastern Iowa Data Center LLC		Premier Petroleum
EverOn Data Center Services	Evocative		Turner Construction Company
Novva Data Centers	Expedient		fifteenfortyseven Critical Systems Realty
Rowan Digital Infrastructure	Five 9s Digital		5Rivers CRE
Digital Core REIT	Freedom Data Centers LLC		Individual owner
US Signal	Involta LLC		NorthPoint Development
EdgeCore Digital Infrastructure	Lifeline Data Centers		Priseda
Edged	Managedway		SDC Capital Partners
NTT Data Group Corporation	Markley Group		TA Realty
CloudHQ	NC Datacenter, LP		Various other single-property owners
Equinix	NV RNO 1 LLC		
American Real Estate Partners (AREP)	Next Tier HD		**Government**
Chirisa Technology Parks	Oso Grande Technologies Inc.		Scott County Board of Supervisors
Element Critical	Rackspace		City of Vernon
Lightedge	Scale Datacenters		
Netrality Data Centers	Skybox Datacenters		
Sabey Data Centers	eVolve Data Center Solutions		
ValorC3 Data Centers	phoenixNAP, LLC		

Note: Owners within each classification are sorted by number of matched properties in the sample, descending.

Table B.IV. Structure of Data Center Mortgage Facilities

	High-OTD lenders		Other lenders	
	Undiv. DC borrowers	Other borrowers	Undiv. DC borrowers	Other borrowers
Share of facility dollars, Post-Jul'21				
CMBS/ABS	0.18	0.32	0.00	0.02
Agent- or trustee-recorded	0.35	0.36	0.00	0.00
Syndicated	0.19	0.05	0.16	0.00
One recorded lender	0.27	0.28	0.84	0.98
Facilities, Post-Jul'21	161	69	73	129
Volume ($ billion), Post-Jul'21	131.5	14.1	8.1	7.5
Volume ($ billion), Pre-Jul'21	3.9	4.9	2.2	2.9

Note: Facilities are new loans and refinancings in the properties' loan histories since January 2015, by lender group and borrower type at origination. Appendix A.3 defines the financing structures.

Appendix C: Debt Measures and Borrower Allocation



Figure C1. Loan Size and Debt Levels Before and After the Boom Onset

Median loan amount per property and debt measures of transactions by lender type in the periods defined in Section III.

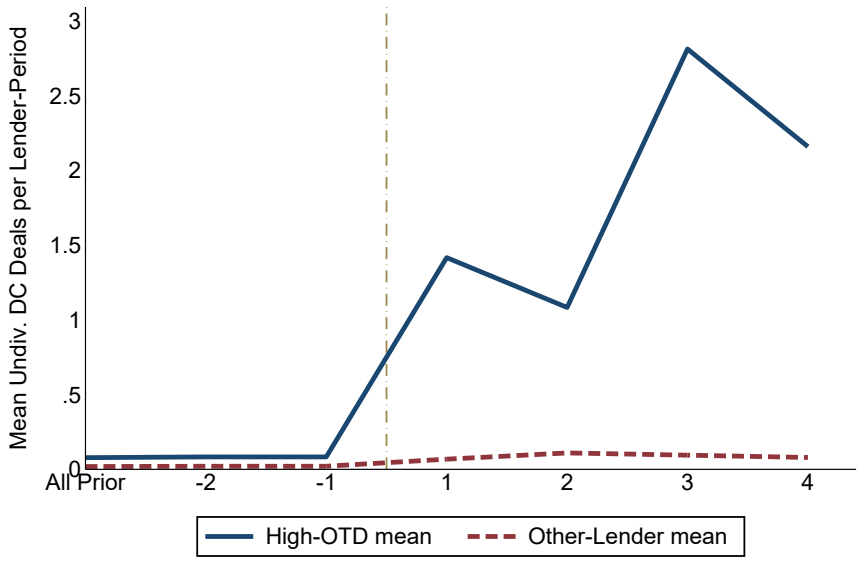


Figure C2. Deals with Undiversified Data Center Borrowers: Levels

Mean Undiversified Data Center deals per lender-period by lender type in the periods defined in Section III.

Table C.I. Pre-Boom Equality Tests

Outcome	Joint F	Joint p	All Prior p	-2 p	Observations	Lender clusters
Loan Amount	0.305	0.737	0.528	0.436	854	266
Debt per Square Foot	0.521	0.595	0.936	0.617	854	266
Debt to Rent Ratio	0.345	0.709	0.872	0.765	854	266
Loan to Value Ratio	0.424	0.655	0.460	0.367	854	266
Undiv. DC Deals	0.002	0.998	0.962	1.000	10,395	385

Note: Each row tests whether the High-OTD gaps in the earliest pre-boom periods (All Prior and period -2) equal the gap in period -1, using the specification of the corresponding event-study figure. The debt-measure and loan-amount rows correspond to Figure 2, and the Undiv. DC Deals row corresponds to Figure 3. The next two columns test each period separately. For the Undiv. DC Deals row, Obs. are lender-periods. Standard errors are clustered at the lender level.

Appendix D: Robustness and Alternative Explanations

Table D.I. Property Characteristics

Characteristic	Pre-July 2021 mean	Coefficient	SE	Observations
Panel A: Real estate characteristics				
RBA (log)	11.314	−0.254	(0.172)	854
Acreage (log)	1.722	−0.240	(0.184)	826
Building age	33.471	0.136	(0.160)	671
Under construction	0.029	−0.036	(0.292)	854
Assessed value per sq ft (USD)	345.6	−0.901	(0.804)	676
CoStar star rating	3.229	−0.233	(0.204)	854
Class A building	0.278	−0.188	(0.233)	854
Renovated	0.290	0.099	(0.138)	854
Panel B: Technology characteristics				
Utility power capacity (log kW)	9.194	−0.017	(0.300)	324
Critical IT capacity (log kW)	8.923	−0.171	(0.512)	167
Utility kW per sq ft	0.177	0.540	(0.796)	324
Power usage effectiveness	1.594	−0.324	(0.985)	127
Data center tier	2.976	0.372	(0.379)	295
Power redundancy	1.754	−0.100	(0.265)	398
Cooling redundancy	1.349	0.152	(0.260)	370

Note: The table reports the High-OTD × Post-July 2021 coefficient from separate regressions for each characteristic. Each characteristic is divided by its standard deviation among pre-July 2021 loans, and the pre-July 2021 mean column is in original units. Panel B covers power capacity, efficiency, and resilience, which *CoStar* reports for part of the sample. Specifications include state and origination-year fixed effects, macroeconomic controls, and local power supply and electricity prices. Standard errors are clustered at the lender level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table D.II. Matching Estimates

	ln(LoanAmount)	ln(DebtPSF)	ln(DebtToRent)	ln(LTV)
High-OTD × Post-Jul'21	1.828***	1.899***	1.899***	2.088***
	(0.509)	(0.466)	(0.461)	(0.385)
Controls and FE	Yes	Yes	Yes	Yes
Observations	630	630	630	630

Note: This table reports estimates of Equation 2, weighted by propensity scores estimated on the controls and fixed effects of Table III, with Low-OTD loans reweighted to resemble High-OTD loans and the sample restricted to common support. The sample is restricted to banks classified as High- or Low-OTD. Controls and fixed effects are as in Table III. Loan amount per property is defined as in Table III, and its column omits RBA from the controls. Standard errors clustered at the lender level are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table D.III. Technology Characteristics and Debt

	(1)	(2)	(3)	(4)
Panel A: ln(*LoanAmount*)				
High-OTD × Post-Jul'21	2.077***	1.974***	1.658***	2.222**
	(0.750)	(0.521)	(0.433)	(1.079)
R^2	0.600	0.562	0.565	0.590
Observations	324	295	337	127
Panel B: ln(*DebtPSF*)				
High-OTD × Post-Jul'21	2.220***	2.026***	1.562***	4.284***
	(0.670)	(0.504)	(0.382)	(1.297)
R^2	0.599	0.512	0.538	0.639
Observations	324	295	337	127
Panel C: ln(*Debt per utility kW*)				
High-OTD × Post-Jul'21	1.985**	1.919	1.441**	4.919
	(0.838)	(1.156)	(0.668)	(3.242)
R^2	0.605	0.760	0.695	0.798
Observations	324	131	163	82
Power capacity and density	Yes			
Data center tier		Yes		
Power and cooling redundancy			Yes	
Power effectiveness (PUE)				Yes
Controls and FE	Yes	Yes	Yes	Yes

Note: Controls and fixed effects are as in Table III. Each column adds the listed *CoStar* characteristics and is estimated on the properties that report them. Without the added characteristics, the estimates on the same properties are similar in size and significance. Debt per utility kW divides debt by *CoStar*'s utility power capacity, and Panel C covers the properties where that capacity is reported. In Panel C, column (1) adds power density only. PUE is power usage effectiveness. Estimates for debt-to-rent and LTV are similar to Panel B. Panel A omits RBA from the controls. Standard errors clustered at the lender level are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table D.IV. Borrower and Lender Fixed Effects

	(1)	(2)	(3)	(4)
Panel A: ln(*LoanAmount*)				
High-OTD × Post-Jul'21	1.298**	1.546**	1.469**	1.863***
	(0.527)	(0.677)	(0.624)	(0.442)
Panel B: ln(*DebtPSF*)				
High-OTD × Post-Jul'21	1.346***	1.540**	1.365**	2.121***
	(0.502)	(0.639)	(0.570)	(0.412)
Observations	520	263	662	569
Borrowers	48	34	150	120
Cells		49	83	
Borrower FE	Yes			
Borrower × event-period FE		Yes		
Lender FE			Yes	
Borrowers with pre-boom borrowing only				Yes
Origination-year FE	Yes		Yes	Yes
Controls	Yes	Yes	Yes	Yes

Note: Controls are as in Table III, and Panel A excludes RBA from the controls, as in Table III. Column 2 uses borrower × event-period fixed effects instead of origination-year fixed effects. Column 4 restricts the sample to borrowers with a mortgage in *CoStar* or a syndicated facility in *DealScan* before July 2021. Observations, Borrowers, and Cells count the properties, borrowers, and fixed-effect groups that identify each estimate. Estimates for debt-to-rent and LTV are similar to Panel B. Standard errors clustered at the lender level are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table D.V. Data Centers and Warehouses

	ln(LoanAmount) (1)	ln(DebtPSF) (2)	ln(DebtToRent) (3)	ln(LTV) (4)
Panel A: Sector-by-state and sector-by-year fixed effects				
High-OTD × Post-Jul'21 × Data Center	1.381**	1.906***	1.949***	1.798***
	(0.684)	(0.565)	(0.539)	(0.609)
High-OTD × Post-Jul'21	0.259	0.100	0.106	0.028
	(0.400)	(0.301)	(0.292)	(0.302)
R^2	0.271	0.167	0.147	0.202
Observations	11,681	11,681	11,476	11,668
Panel B: Adds lender-by-year fixed effects				
High-OTD × Post-Jul'21 × Data Center	2.221***	2.261***	2.171***	1.929***
	(0.725)	(0.547)	(0.512)	(0.530)
R^2	0.671	0.565	0.548	0.591
Observations	8,340	8,340	8,171	8,328
Controls and FE	Yes	Yes	Yes	Yes

Note: This table reports estimates of Equation 4 for data center and warehouse loans, with the controls of Table III. Column 1 omits RBA from the controls. Outcomes and lender types are defined the same way in both sectors. The second row in Panel A is the post-boom change at High-OTD lenders on warehouses. Panel B adds lender-by-year fixed effects. Standard errors clustered at the lender level are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table D.VI. Bank Size

	ln(LoanAmount)	ln(DebtPSF)	ln(DebtToRent)	ln(LTV)	Undiv. DC Deals
Panel A: Predetermined bank-assets control, bank lenders only					
High-OTD × Post-Jul'21	2.172***	2.086***	2.061***	2.351***	2.141**
	(0.488)	(0.477)	(0.489)	(0.504)	(0.936)
ln(Bank assets) × Post-Jul'21	−0.012	−0.008	−0.008	−0.019	0.086
	(0.079)	(0.075)	(0.073)	(0.082)	(0.083)
Bank-size controls	Yes	Yes	Yes	Yes	Yes
Observations / lender-periods	650	650	650	650	3,408
Panel B: Excluding the largest lenders					
5 largest					
High-OTD × Post-Jul'21	1.330***	1.536***	1.509***	1.368***	1.649**
	(0.463)	(0.422)	(0.393)	(0.421)	(0.670)
Observations / lender-periods	771	771	771	771	10,233
10 largest					
High-OTD × Post-Jul'21	0.952*	1.139**	1.245***	1.032**	1.466**
	(0.545)	(0.448)	(0.448)	(0.501)	(0.732)
Observations / lender-periods	676	676	676	676	10,071
15 largest					
High-OTD × Post-Jul'21	1.077*	1.298***	1.388***	1.041**	1.051*
	(0.563)	(0.448)	(0.449)	(0.518)	(0.599)
Observations / lender-periods	594	594	594	594	9,936
Controls and FE	Yes	Yes	Yes	Yes	Yes

Note: The debt-measure columns use the controls and fixed effects of Table III, and the Undiv. DC Deals column uses the design of Table IV. The bank-size controls in Panel A are defined in Appendix A.3. Panel B excludes the 5, 10, or 15 largest bank lenders by 2020 total assets. Column 1 omits RBA from the controls. Standard errors clustered at the lender level are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table D.VII. OTD Classification

	ln(LoanAmount)	ln(DebtPSF)	ln(DebtToRent)	ln(LTV)	Undiv. DC Deals
Panel A: 2022–2025 average classification					
High-OTD × Post-Jul'21	1.543***	1.693***	1.720***	1.586***	1.835***
	(0.392)	(0.355)	(0.340)	(0.372)	(0.648)
Observations / lender-periods	854	854	854	854	10,395
Panel B: Excluding trustee- and agent-recorded loans					
High-OTD × Post-Jul'21	1.738***	1.794***	1.807***	1.832***	1.345**
	(0.398)	(0.378)	(0.363)	(0.367)	(0.570)
Observations / lender-periods	731	731	731	731	10,260
Panel C: Excluding foreign banks					
High-OTD × Post-Jul'21	1.641***	1.839***	1.806***	1.540***	1.573**
	(0.407)	(0.354)	(0.339)	(0.370)	(0.640)
Observations / lender-periods	701	701	701	701	9,909
Panel D: Continuous OTD score					
OTD score × Post-Jul'21	0.120**	0.127**	0.134**	0.128**	0.657***
	(0.052)	(0.057)	(0.053)	(0.057)	(0.138)
Observations / lender-periods	465	465	465	465	2,664
Panel E: Binary High-OTD benchmark					
High-OTD × Post-Jul'21	2.063***	2.116***	2.046***	2.138***	2.993**
	(0.451)	(0.445)	(0.446)	(0.413)	(1.286)
Observations / lender-periods	465	465	465	465	2,664
Controls and FE	Yes	Yes	Yes	Yes	Yes

Note: This table re-estimates Equation 2 and the borrower-allocation design of Table IV under alternative High-OTD classifications. Loan amount is defined as in Table III, and its column omits RBA from the controls. Panel A uses 2022–2025 *Call Report* data. Panel B excludes loans recorded under a trustee, servicer, or loan agent. Panel C excludes foreign banks. Panel D uses the standardized 2017–2020 average OTD score for banks classified from their own *Call Reports*, and Panel E reports the binary specification on the same banks. Standard errors clustered at the lender level are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table D.VIII. Cutoff Date

Cutoff date	ln(LoanAmount)	ln(DebtPSF)	ln(DebtToRent)	ln(LTV)	Undiv. DC Deals
Jan 2021 (−6 months)	1.703***	1.711***	1.694***	1.763***	1.467***
	(0.388)	(0.361)	(0.363)	(0.375)	(0.532)
Apr 2021 (−3 months)	1.740***	1.858***	1.858***	1.779***	1.583***
	(0.377)	(0.353)	(0.343)	(0.347)	(0.571)
Oct 2021 (+3 months)	1.396***	1.453***	1.481***	1.366***	1.692***
	(0.396)	(0.371)	(0.356)	(0.393)	(0.628)
Jan 2022 (+6 months)	1.206***	1.282***	1.335***	1.193***	1.757***
	(0.426)	(0.380)	(0.364)	(0.415)	(0.644)
Controls and FE	Yes	Yes	Yes	Yes	Yes
Observations / lender-periods	854	854	854	854	10,010–10,395

Note: The first four columns re-estimate Equation 2 with the Post-Jul'21 cutoff shifted three and six months earlier and later. The Undiv. DC Deals column re-estimates Table IV's design with the 15-month periods rebuilt around each shifted cutoff. The Jul 2021 row reproduces Table IV. Column 1 omits RBA from the controls. Standard errors clustered at the lender level are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table D.IX. Sample and Refinancing Robustness

	ln(LoanAmount)	ln(DebtPSF)	ln(DebtToRent)	ln(LTV)	Undiv. DC Deals
Panel A: Single-property loans					
High-OTD × Post-Jul'21	1.556***	1.481***	1.470***	1.314***	0.390**
	(0.502)	(0.467)	(0.446)	(0.474)	(0.155)
Observations / lender-periods	448	448	448	448	8,748
Panel B: Excluding non-operational properties					
High-OTD × Post-Jul'21	1.369***	1.410***	1.402***	1.320***	1.451***
	(0.401)	(0.317)	(0.307)	(0.351)	(0.526)
Observations / lender-periods	637	637	637	637	8,640
Panel C: Excluding refinancings					
High-OTD × Post-Jul'21	1.362**	1.481***	1.491***	1.353***	1.288***
	(0.556)	(0.520)	(0.513)	(0.507)	(0.464)
Observations / lender-periods	631	631	631	631	10,395
Panel D: Controlling for refinancing					
High-OTD × Post-Jul'21	1.589***	1.747***	1.755***	1.648***	1.720***
	(0.397)	(0.355)	(0.343)	(0.374)	(0.611)
Observations / lender-periods	854	854	854	854	10,395
Controls and FE	Yes	Yes	Yes	Yes	Yes

Note: This table reports estimates of Equation 2. Panel A restricts to single-property loans. Panel B drops properties classified by *CoStar* as abandoned, proposed, final planning, or under construction. Panel C excludes loans classified as refinancings. Panel D controls for refinancing, with a refinancing indicator in the debt-measure columns and each lender's refinancing share interacted with Post-Jul'21 in the Undiv. DC Deals column. The debt-measure columns use the controls and fixed effects of Table III, and the Undiv. DC Deals column uses the design of Table IV. Column 1 omits RBA from the controls. Standard errors clustered at the lender level are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table D.X. Borrower Allocation: Alternative Measures

	Most recent loan	Loans	Active-period zero-filling	PPML
High-OTD × Post-Jul'21	1.543***	0.710***	4.656***	1.646***
	(0.543)	(0.263)	(1.649)	(0.394)
Controls and FE	Yes	Yes	Yes	Yes
Observations / Lender-periods	7,182	10,287	736	5,103

Note: All columns use the design of the first column of Table IV. Most recent loan counts only each property's most recent loan, with the current owner's borrower type. Loans counts loans rather than financed properties, so a loan secured by several properties counts once. Active-period zero-filling counts zero deals only between a lender's first and last observed loan. PPML is a Poisson model with lender and period fixed effects and period length as exposure, estimated on lenders with at least one Undiversified Data Center deal. Standard errors clustered at the lender level are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.