Datacenter Mortgages and the Originate-to-Distribute Model*

Pedro Gete[†] Amparo Mercader[‡]

June 2026

Abstract

The AI boom has fueled a rapid expansion of datacenter infrastructure financed by banks that distribute rather than retain credit risk. We ask whether originate-to-distribute incentives shape underwriting in this market. Banks with stronger distribution orientation originate datacenter loans with leverage ratios roughly 3 to 4 times as high as other lenders and are more likely to finance borrowers whose revenues depend entirely on datacenter demand. These patterns are consistent with originate-to-distribute moral hazard operating in a setting where collateral is specialized, outside investors face large information asymmetries, and no prior default cycle provides risk benchmarks.

Keywords: Artificial Intelligence, Datacenter, Leverage, Credit Risk, Originate-to-Distribute, Digital Infrastructure

JEL Classification: G12, G21, G23, R31.

*We thank Isabel de Valenzuela, Carlos Garriga, Lucas Gulão, Stijn Van Nieuwerburgh and participants at the 2026 Dallas Fed Conference and the University of Alicante Housing and Macroeconomy Workshop for helpful comments. This paper was partially funded by MICIU/AEI/10.13039/501100011033/FEDER, UE, Grant No. PID2024-161318NB-I00.

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

[‡]IE University and Georgetown University. Email: mercada@georgetown.edu

1 Introduction

Datacenter mortgage originations exceeded $110 billion over the past four years, with activity accelerating sharply after 2021. This expansion has been financed by a concentrated set of large institutions that are also active distributors of credit risk through syndication, securitization, and secondary loan sales. We ask whether the incentive to distribute rather than retain credit shapes underwriting outcomes in this market.

When a lender expects to sell or syndicate a loan, the quality of the borrower has a smaller effect on the lender's future income than when the loan is retained. Screening effort may be weakened, and originated loans may be of lower average quality than outside investors anticipate. This originate-to-distribute mechanism has been documented in residential mortgage markets, subprime consumer credit, and corporate syndicated lending (Bord and Santos, 2015; Dell'Ariccia et al., 2012; Keys et al., 2010; Purnanandam, 2011). Datacenter lending provides a new setting in which to examine whether the same forces operate: the market is growing rapidly and is dominated by distribution-oriented institutions.

We construct a novel property-loan-level dataset combining commercial real estate records from CoStar, covering approximately 1,850 datacenter properties in the United States as of November 2025, with mortgage origination data and lender characteristics from quarterly Call Report data accessed via the Federal Reserve Bank of New York. Following Purnanandam (2011), we construct bank-level originate-to-distribute scores from the ratios of loan assets, trading assets, and repurchase agreements to total assets, classify lenders into High-OTD and Low-OTD groups, and compare lending outcomes across lender types for otherwise comparable properties.

Properties financed by High-OTD lenders carry debt per square foot, debt-to-rent ratios, and loan-to-value ratios that are approximately 3.4, 3.6, and 3.4 times as high as those financed by other lenders, conditional on property size, local macro conditions, tenancy type, taxes per square foot, state fixed effects, and loan-origination year fixed effects. These patterns survive matched comparisons enforcing overlap on observables, exclusion of the largest banks, and instrumental variable estimation using predetermined 2021 balance-sheet measures.

A one-standard-deviation increase in the continuous OTD score predicts a 6.0 percentage point higher probability of financing an Undiversified Datacenter owner, a borrower whose revenues are concentrated in datacenter operations. The OTD leverage premium is driven entirely by loans to these specialized landlords: among diversified owners the OTD gradient is indistinguishable from zero, while among Undiversified Datacenter owners it is large and significant. Distribution-oriented lenders originate most aggressively precisely when financing the riskiest borrowers.

This paper extends the originate-to-distribute evidence to digital infrastructure finance and shows that OTD moral hazard is amplified in emerging asset classes where information asymmetries are unusually severe (DeMarzo and Duffie, 1999; Gorton and Pennacchi, 1995). Unlike traditional real estate, datacenter values depend on power availability, cooling capacity, network connectivity, and tenant contracts that outside investors cannot easily observe or verify. The party originating the loan possesses substantially more information about collateral quality than the investors ultimately bearing the risk, and no prior default cycle has established the benchmarks that discipline aggressive origination in established markets. The patterns of aggressive underwriting and borrower selection documented in residential mortgages, corporate syndicated lending, and commercial real estate securitization (Black et al., 2012; Bord and Santos, 2015; Purnanandam, 2011) are present here in a setting where the conditions for moral hazard are more extreme. The concentration of high-leverage origination among distribution-oriented lenders in this environment may be of interest to supervisors and investors monitoring the financing of AI infrastructure.

Our identification strategy exploits cross-bank variation in originate-to-distribute orientation measured from predetermined balance-sheet characteristics. Following Purnanandam (2011), we construct an OTD score from banks' holdings of loans, trading assets, and repurchase agreements and classify lenders according to their propensity to distribute rather than retain credit risk. To address concerns that OTD orientation captures bank size, funding advantages, or conditions specific to the post-2021 datacenter boom, we instrument for lender OTD orientation using balance-sheet characteristics measured in 2021, before the acceleration in datacenter lending and before the monetary tightening cycle. The instruments strongly predict subsequent OTD status and generate estimates close to baseline OLS. A pre-boom placebo shows that the same instruments have little explanatory power before the datacenter boom, supporting the interpretation that OTD incentives became relevant as the market scaled. Matching estimators, sensitivity analysis following Oster (2019), and bank-size controls further confirm that the results are not driven by observable or unobservable differences in borrower or property characteristics.

The remainder of the paper proceeds as follows. Section 2 reviews the related literature. Section 3 describes the data. Section 4 documents recent dynamics in datacenter mortgage debt. Section 5 describes the OTD model in the datacenter setting. Section 6 presents the empirical analysis and identification strategy. Section 7 concludes.

2 Related Literature

When a bank expects to sell or syndicate a loan rather than retain it, the quality of the borrower has a smaller effect on the lender's future income. Screening effort may therefore fall and originated loans may be of lower average quality than outside investors anticipate (DeMarzo and Duffie, 1999; Drucker and Puri, 2009; Gorton and Pennacchi, 1995). Reputational concerns and retained exposure can partially offset these incentives (Puri et al., 2011; Sufi, 2007), but theory predicts both are weakest in novel asset classes without established benchmarks and in fast-growing markets where origination fees are large relative to reputational stakes.

The empirical evidence on originate-to-distribute incentives spans residential mortgages, consumer credit, and corporate lending. Keys et al. (2010) show that securitization reduced screening effort in residential mortgages, and Purnanandam (2011) shows that banks with higher originate-to-distribute propensity, measured from regulatory balance sheets, originated loans with higher ex post default rates. Dell'Ariccia et al. (2012) document a broader deterioration of lending standards in high-securitization areas. In corporate credit, Bord and Santos (2015) and Ivashina and Sun (2011) show that arrangers with higher distribution propensity impose looser covenants and select riskier borrowers, and Gustafson et al. (2021) find that active monitoring declines with loan maturity. Originate-to-distribute propensity is measured using balance-sheet ratios following Purnanandam (2011); McGowan et al. (2024) and McGowan and Nguyen (2023) provide recent extensions.

In commercial real estate, Black et al. (2012) study over 30,000 commercial mortgage-backed securities loans spanning office, retail, industrial, and multifamily properties and find that conduit lenders, institutions that originate loans for distribution rather than balance-sheet retention, produce loans with significantly higher delinquency rates than balance-sheet lenders. The result extends originate-to-distribute incentive problems beyond residential and corporate credit to secured commercial lending. We study the same mechanism in digital infrastructure, a sector where these incentives have not previously been examined.

The academic literature on datacenter finance is nascent. Darmouni et al. (2025) study the energy cost implications of AI and datacenter investment, and the financing structures supporting the current buildout remain largely unstudied. We provide evidence on how lender incentives shape origination in a market that is large, growing rapidly, and concentrated among a small set of distribution-oriented institutions.

3 Database

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 November 2025 (approximately 1,850 assets). From this universe, we identify 550 properties with mortgage records, which form the mortgaged sample used in the leverage analysis. For each property, we observe rentable building area (RBA), geographic location, development status, shell rent, property taxes, percent leased, year built, ownership information, and loan origination records. CoStar's loan records include origination amount, origination date, interest rate, loan type, and originating lender name.

We construct property-level leverage measures from these 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.

Owner taxonomy. We classify datacenter owners into four categories based on business scope and revenue diversification. Undiversified Datacenters are firms whose primary business is owning and operating datacenter real estate, with revenues concentrated in datacenter leasing and related services. Big Tech owners are large technology firms with substantial non-datacenter revenue streams. Government owners are public-sector entities. The residual category, Other, includes diversified private owners such as industrial REITs, institutional real estate investors, and other corporations with broader revenue bases. The key distinction for our analysis is between specialized datacenter landlords and diversified owners, as this affects both the borrower's risk profile and the lender's ability to assess creditworthiness. Panel C of Table 1 reports median debt per square foot by owner type: Undiversified Datacenter owners carry roughly 2.7 times as much debt per square foot as diversified owners.

Lender taxonomy. We link lenders to quarterly Call Report data (FFIEC), accessed via the Federal Reserve Bank of New York data products covering 2022 to 2025. Following

Purnanandam (2011), we construct a composite OTD score for each bank from three balance-sheet ratios: loan assets to total assets (a proxy for retention), and trading assets and repurchase agreements to total assets (proxies for distribution-oriented intermediation). Each ratio is standardized within year and the score is defined as

$$\text{OTD Score} = -z(\text{Loan Assets}) + z(\text{Trading Assets}) + z(\text{Repo Assets}). \quad (1)$$

Banks are ranked by their average score over 2022 to 2025 and those in the top two deciles are classified as High-OTD. Appendix Table B1 lists the full lender assignments and earnings-call transcripts validate the classifications. Appendix Table B2 reports group means: High-OTD banks hold 40.2% of assets in loans versus 69.7% for Low-OTD banks, and hold substantially more in trading assets and repo positions. The remaining observations include non-bank lenders (insurance companies, private credit funds) and government and cooperative lenders.

As an external validation, we match our lenders to DealScan syndicated loan data covering datacenter-related borrowers over 2019 to 2025. High-OTD banks act as lead arrangers in 65.6% of syndicated datacenter deals versus 47.8% for Low-OTD banks, and participate in Term Loan B structures at nearly twice the rate, confirming that the balance-sheet classification identifies the same institutions as distribution-oriented in their broader lending activity (Appendix Table B3).

4 Recent Dynamics of Datacenter Mortgage Debt

This section documents the recent expansion of U.S. datacenters and the mortgage debt associated with them. Figure 1 maps total datacenter capacity at the county level. Capacity is concentrated in Northern Virginia and a small number of additional metropolitan hubs, reflecting clustering around power infrastructure and network exchange points.

Figure 2 reports cumulative origination of datacenter mortgage debt by lender type. Originations accelerate sharply after 2021, with High-OTD banks accounting for the dominant share of this expansion. Before 2021 all lender types originate at similar and modest rates. After 2021 the High-OTD bank line rises nearly vertically, reaching approximately $95 billion in cumulative originations by 2025, while Low-OTD banks, non-bank lenders, and government lenders remain flat by comparison. This pattern indicates that the post-2021 datacenter lending boom was financed predominantly by distribution-oriented institutions, establishing the conditions under which originate-to-distribute incentives are most likely to operate.

Figure 3 shows that the acceleration reflects not only more loans but substantially larger ones: median loan size rises sharply after 2021 and the interquartile range widens considerably, consistent with the emergence of large portfolio financings alongside the broader lending boom.

Figure 4 documents the distributional shift in more detail. The upper tail of the loan size distribution thickens considerably after 2021, with aggregate debt increasingly concentrated in very large loans including financings in the $3 to $7 billion range. These mega-loans represent a shift toward portfolio-level financing structures that collateralize multiple datacenter properties simultaneously, raising questions about the extent to which individual property fundamentals are assessed when loans are originated at the portfolio level. Together, these patterns establish the conditions under which OTD incentives are most likely to operate: rapid origination growth, rising loan sizes, and a concentrated set of lenders.

5 The OTD Model Then and Now

The 2008 crisis established the core empirical pattern: banks organized around origination and distribution weakened screening as volume grew, passing risk to investors who could not observe loan quality. The datacenter boom reproduces this structure, with one important difference. Residential mortgages were backed by assets whose values, however inflated, were widely observable. Datacenter mortgages are backed by specialized infrastructure whose value depends on power capacity, cooling, and tenant contracts, creating information frictions that increase the scope for aggressive origination when distribution incentives are strong.

The combination of distribution-oriented lenders, specialized collateral, and the absence of a prior default cycle creates conditions under which OTD moral hazard is likely to be most severe. The mechanisms that discipline aggressive origination in established markets, namely reputational benchmarks, investor experience, and observable collateral values, are all absent simultaneously.

The empirical analysis tests whether High-OTD banks originate more aggressively in this setting and whether OTD orientation predicts borrower specialization. We address alternative explanations including bank size, funding costs, borrower sorting, property quality, unobservable selection, and macroeconomic timing through observable controls and matching, bank-size robustness checks, sensitivity to unobservables, and instrumental variables using predetermined 2021 balance-sheet measures validated by a pre-boom placebo.

6 Empirical Analysis

The central identification challenge is distinguishing OTD moral hazard from six alternative explanations. First, large banks may lend more aggressively simply because of greater capacity and broader borrower relationships, not because of distribution incentives. Second, banks with cheaper or more abundant funding may originate larger loans regardless of OTD orientation. Third, riskier borrowers may systematically sort toward High-OTD banks rather than being selected by them. Fourth, High-OTD banks may finance intrinsically larger or higher-quality properties that warrant more debt. Fifth, unobserved asset characteristics correlated with both lender type and leverage could explain the premium without any moral hazard. Sixth, the post-2021 boom may have coincided with both more OTD activity and more aggressive lending for unrelated cyclical reasons. We address each of these alternatives through observable controls and matching, bank-size robustness checks, sensitivity to unobservables, and instrumental variables.

6.1 Baseline Results

Our baseline specification compares leverage outcomes at properties financed by High-OTD banks to all other properties:

$$Y_i = \beta_H \,\text{HighOTD}_i + \mathbf{X}_i'\theta + \gamma_{s(i)} + \delta_{t(i)} + \varepsilon_i, \tag{2}$$

where HighOTD_i equals one if property i is financed by a High-OTD bank. The control vector $\mathbf{X}_i$ includes ln(RBA), local macro controls, tenancy type, and taxes per square foot. All specifications include state fixed effects $\gamma_{s(i)}$ and loan-origination year fixed effects $\delta_{t(i)}$, with standard errors clustered at the lender level. We estimate this specification using three dependent variables: ln(DebtPSF), ln(DebtToRent), and ln(LTV).

Table 2 reports the results. High-OTD banks originate loans with substantially higher leverage across all three outcomes: properties they finance carry debt per square foot, debt-to-rent, and loan-to-value ratios approximately 3.4, 3.6, and 3.4 times as high as those financed by other lenders. Table 3 adds a Low-OTD indicator to the same specification, with non-bank and government lenders as the omitted category. The High-OTD coefficient rises to 1.60 because the omitted group now excludes Low-OTD banks from the baseline. The leverage premium remains concentrated among High-OTD banks: the Low-OTD coefficient is statistically indistinguishable from zero, and Wald tests reject equality of the two coefficients with p-values of 0.0005, 0.0001, and 0.0007 across the three outcomes.

The standalone Low-OTD results are reported in Appendix Table C1; the negative sign there reflects the broader omitted category, which includes non-bank and government lenders alongside High-OTD banks. The results are unchanged when restricting to loans collateralized by a single property, ruling out RBA allocation mechanics as a driver (Appendix Table C2).

6.2 Borrower Selection

We test whether OTD orientation also predicts the type of borrower financed. Specifically, we estimate whether lenders with higher OTD scores are more likely to finance Undiversified Datacenter owners, whose revenues are concentrated in datacenter operations and who present greater credit risk than diversified owners:

$$\text{UndiversifiedDatacenter}_i = \beta\,\text{OTDz}_i + \mathbf{X}_i'\theta + \gamma_{s(i)} + \delta_{t(i)} + \varepsilon_i, \tag{3}$$

where OTDz_i is the lender's standardized OTD score and standard errors are clustered at the lender level.

Table 4 reports the results. A one-standard-deviation increase in OTD score predicts a 6.0 percentage point higher probability of financing an Undiversified Datacenter owner, stable across specifications with and without controls and fixed effects. Higher OTD propensity therefore predicts not only greater leverage but selection into borrowers with more concentrated datacenter exposure. This result also survives matched comparisons: within samples of comparable properties, High-OTD banks remain more likely to finance Undiversified Datacenter owners, with a matched estimate of 14 percentage points under propensity-score weighting.

6.3 Borrower Type and the OTD Leverage Premium

The leverage and borrower-selection results raise a natural question: does the OTD leverage premium vary by borrower type? Figure 5 provides a first answer. The left panel shows that the OTD leverage premium is positive and monotonic across the full sample. The right panel shows that the premium is concentrated among Undiversified Datacenter owners and negligible among diversified owners, consistent with distribution-oriented lenders originating most aggressively when financing specialized datacenter landlords.

We test this formally by interacting the continuous OTD score with an indicator for

Undiversified Datacenter owners:

$$\begin{aligned} Y_i = {} & \beta \, \mathrm{OTDz}_i + \lambda \, \mathrm{UndiversifiedDC}_i + \phi \, (\mathrm{OTDz}_i \times \mathrm{UndiversifiedDC}_i) \\ & + \mathbf{X}_i' \theta + \gamma_{s(i)} + \delta_{t(i)} + \varepsilon_i. \end{aligned} \tag{4}$$

Appendix Table C3 reports the results. The interaction term is positive across all three leverage outcomes and significant at the 10 percent level. Among Undiversified Datacenter owners, a one-standard-deviation increase in OTD score is associated with roughly 28 to 32 percent higher leverage across the three measures.

6.4 Identification

We next address alternative explanations for the leverage premium and borrower-selection patterns. The baseline OLS specification already controls for property quality and macroeconomic timing through rich controls and fixed effects. We further show the results are not driven by observable sorting, bank size, or unobservable selection, and instrument for High-OTD status using predetermined 2021 balance-sheet measures that vary by bank business model before the datacenter lending boom and before the March 2022 rate-hike cycle.

6.4.1 Observable Controls and Matching

To address observable sorting, we restrict the sample to High-OTD and Low-OTD banks and enforce overlap in observables. We estimate the probability that a loan is financed by a High-OTD bank using a probit model with the baseline controls and fixed effects, restrict to the region of common support, and estimate the High-Low difference using inverse probability weighting. As a complementary approach, we implement coarsened exact matching by forming strata based on state and property-size quartiles. Table 5 reports the results. Across both approaches the High-OTD premium remains positive and precisely estimated. The matched estimates of 0.76 to 0.83 are lower than the full-sample OLS coefficients, which is expected since the matched sample restricts to bank lenders only and enforces overlap on observables, both of which reduce the comparison group relative to the baseline specification. The estimates are nonetheless economically large, representing leverage ratios roughly twice as high among High-OTD bank loans within samples of comparable properties.

6.4.2 Bank Size

All 16 banks with total assets exceeding $100 billion as of 2021 Q4 are classified as High-OTD, including JPMorgan Chase, Bank of America, Wells Fargo, and Goldman Sachs. We address the concern that the leverage premium reflects balance-sheet capacity rather than distribution incentives in two ways. Excluding these 16 banks entirely, Table 6 shows that $\hat{\beta}_H$ remains large and precisely estimated across all three leverage outcomes, with coefficients larger than the full-sample baseline. Augmenting the baseline with each lender's 2021 Q4 total assets as a continuous control, the bank asset coefficient is economically negligible and statistically indistinguishable from zero while $\hat{\beta}_H$ remains positive and precisely estimated (Appendix Table C4).

6.4.3 Selection on Unobservables

Sensitivity analysis following Oster (2019) confirms that unobservable property characteristics would need to be more than four times stronger than all observable controls combined to reduce the leverage coefficients to zero, and nearly three times stronger to explain the loan size gap between High-OTD and Low-OTD banks (Appendix Table C5).

6.4.4 Instrumental Variables

The predetermined 2021 balance-sheet instruments provide the cleanest separation of OTD orientation from bank size, funding conditions, and macroeconomic timing. The instruments vary by bank business model as of 2021 and are therefore unrelated to the post-2021 datacenter boom. They strongly predict High-OTD status, with a Kleibergen-Paap

F-statistic of 47.70, well above the Stock-Yogo 10% critical value of 19.93.[1]

Table 7 reports the 2SLS estimates: 1.258 for ln(DebtPSF) and 1.229 for ln(LTV), stable across balance-sheet and income-based instrument sets. The 2SLS estimates exceed OLS by approximately 60%, consistent with classical measurement error in the binary High-OTD classification attenuating OLS toward zero.

Applying the same instruments to the 2015–2019 pre-boom sample yields a Kleibergen-Paap F-statistic of 4.01 and statistically insignificant second-stage estimates, while the 2022–2025 baseline produces a strong first stage ($F = 47.70$) and significant estimates (Appendix Table D2). The OTD channel operates during the post-2021 boom, not before it.

Applying the instruments to the borrower selection equation, the 2SLS estimate of 5.1 percentage points is significant at the 5 percent level and close to the OLS estimate of 6.0 percentage points (Appendix Table D3), consistent with the same predetermined balance-sheet characteristics predicting both aggressive leverage and selection into Undiversified Datacenter owners.

Two further checks validate the IV strategy. Adding the brokered deposit ratio as a control, the deposit coefficient is statistically insignificant while the High-OTD estimates remain economically large, ruling out a funding cost interpretation (Appendix Table D4). The 2SLS estimates are stable across alternative instrument sets including loan-to-assets only, trading-assets only, and noninterest income share, confirming the results are not sensitive to the choice of excluded instruments (Appendix Table D5).

7 Conclusion

This paper studies whether originate-to-distribute incentives shape underwriting in datacenter mortgage markets. High-OTD banks finance debt per square foot, debt-to-rent, and loan-to-value ratios roughly 3.4, 3.6, and 3.4 times as high as other lenders, and a one-

[1] The instruments are highly persistent across years, with cross-bank correlations ranging from 0.867 to 0.995 over 2022–2025; full persistence correlations and first-stage diagnostics are reported in Appendix Table D1.

standard-deviation increase in the continuous OTD score predicts a 6.0 percentage point higher probability of financing an Undiversified Datacenter owner. The OTD Leverage Premium is concentrated among Undiversified Datacenter owners, the borrowers with the most concentrated datacenter and AI exposure.

These patterns survive matched comparisons enforcing overlap on observables and exclusion of the largest banks. Instrumental variable estimates using predetermined 2021 balance-sheet measures are close to OLS and stable across instrument sets, and a pre-boom placebo confirms the OTD channel is relevant only during the post-2021 expansion. Together these findings extend the originate-to-distribute literature to digital infrastructure finance, a setting where distribution incentives interact with unusually severe information asymmetries, specialized collateral, and the absence of any prior default cycle. The evidence is consistent with OTD moral hazard being amplified precisely in the conditions that characterize emerging infrastructure markets. We do not directly observe ex-post loan performance, and the interpretation of the leverage and borrower selection differentials as reflecting weaker screening depends on the identifying assumptions discussed in Section 6.

Several questions remain for future research. More granular data on loan-level distribution decisions would allow a sharper test of the originate-to-distribute mechanism than balance-sheet proxies permit. Linking origination outcomes to subsequent loan performance would establish whether the leverage and borrower selection differentials documented here translate into realized differences in credit quality. The welfare implications depend on who ultimately bears the distributed risk and how effectively those investors monitor the underlying collateral, questions that the geographic concentration of datacenter capacity and the novelty of the asset class make difficult to answer at this stage.

References

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.

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.

Drucker, S. and Puri, M. (2009). On loan sales, loan contracting, and lending relationships. *The Review of Financial Studies*, 22(7):2635–2672.

Ghent, A. C. and Kudlyak, M. (2011). Recourse and residential mortgage default: Evidence from u.s. states. *The Review of Financial Studies*, 24(9):3139–3186.

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.

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*. Early View / Epub ahead of print (online: 2024-10-01).

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.

Oster, E. (2019). Unobservable selection and coefficient stability: Theory and evidence. *Journal of Business & Economic Statistics*, 37(2):187–204.

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.

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

Figures

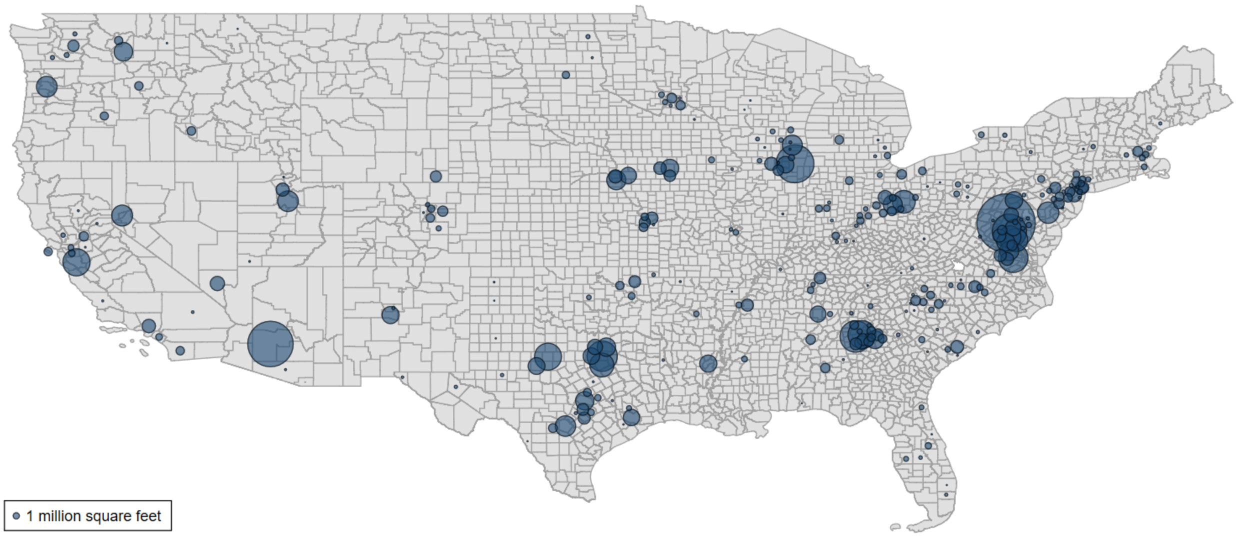

Figure 1. Datacenter Capacity by County

This figure maps county-level datacenter capacity using CoStar inventory data as of November 2025. Capacity is measured as total rentable building area across existing and under-construction properties. Marker size is proportional to county-level capacity.

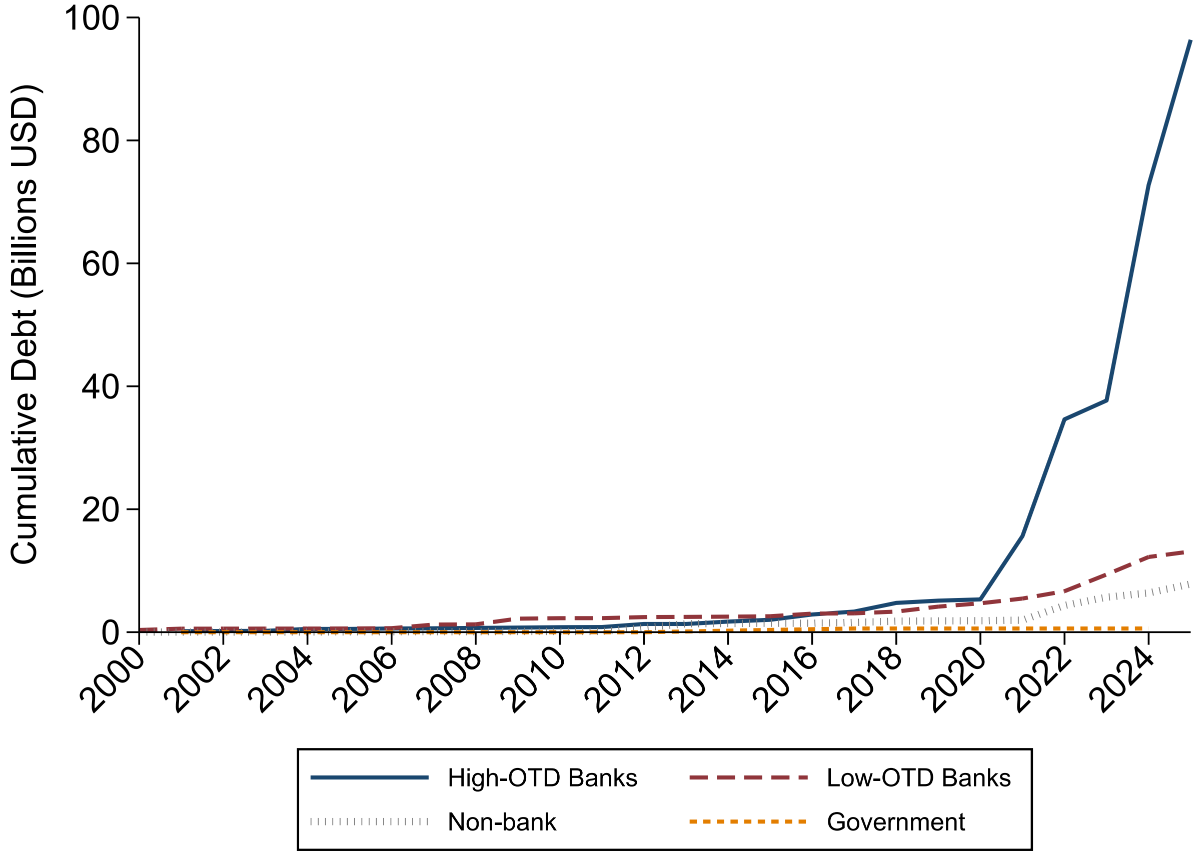

Figure 2. Cumulative Datacenter Mortgage Origination by Lender Type

This figure plots cumulative datacenter mortgage origination by lender type. High-OTD banks are institutions in the top two deciles of the balance-sheet OTD score defined in Equation 1. Loan amounts are summed cumulatively over origination year.

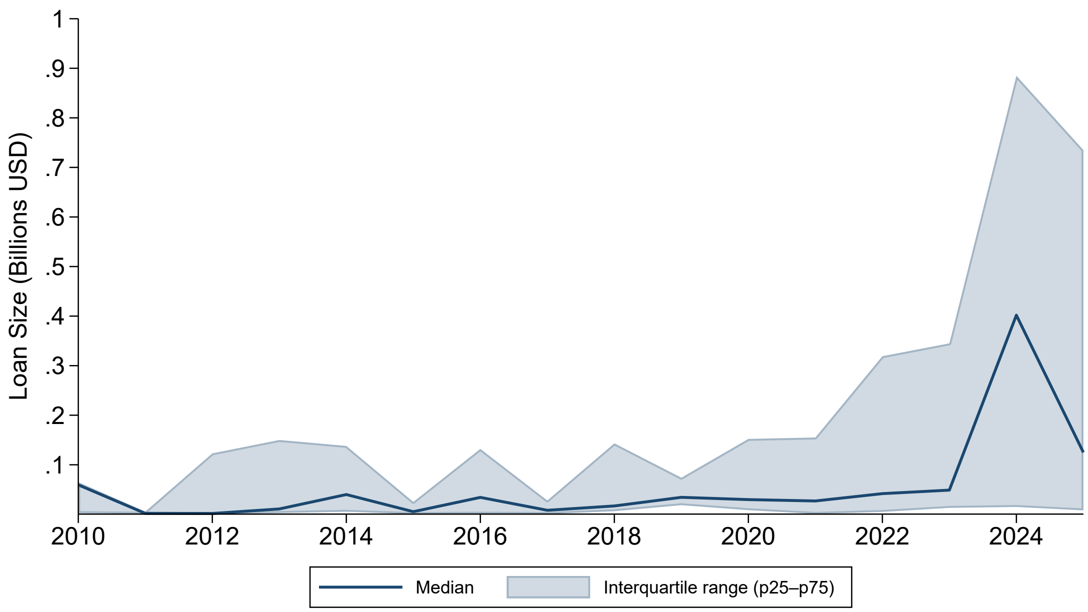

Figure 3. Datacenter Loan Size over Time

This figure plots the annual distribution of property-level datacenter loan sizes. The solid line reports the median loan size by origination year, while the shaded band shows the interquartile range (25th to 75th percentiles). Loan sizes are measured in billions of U.S. dollars and are constructed at the loan level, with each loan counted once in its origination year.

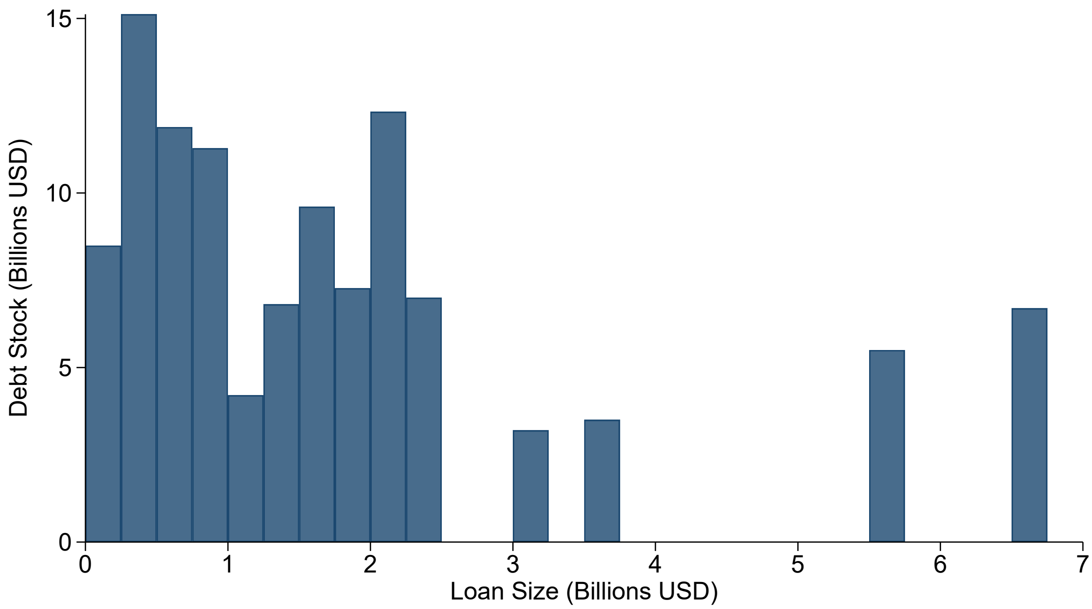

Figure 4. Aggregate Datacenter Debt by Loan Size

This figure plots the distribution of aggregate datacenter debt across loan-size bins. Each bar shows the total outstanding debt (in billions of U.S. dollars) within $0.25 billion loan-size bins, aggregated across all loans originated between 2010 and 2025.

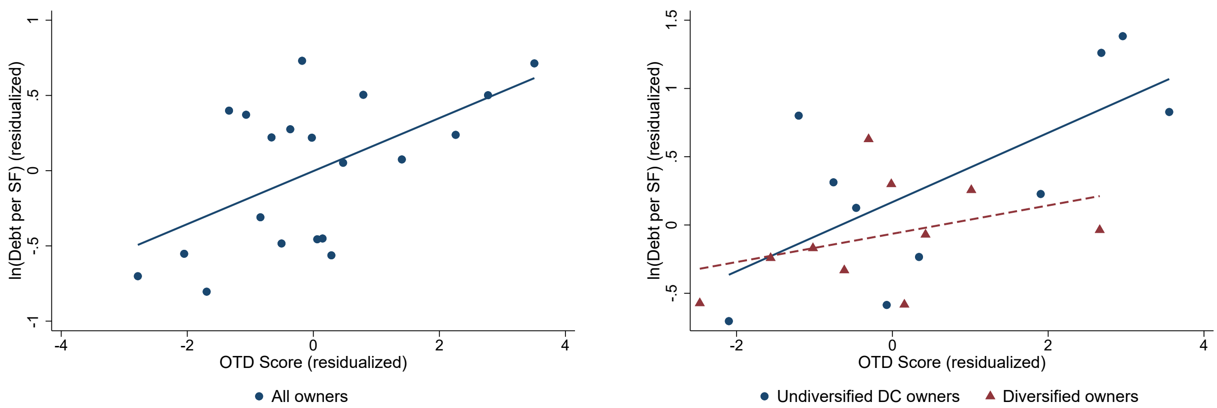

Figure 5. OTD Orientation and Datacenter Leverage

Each point represents a decile of the OTD score distribution. Both axes are residualized on property size, local macro conditions, tenancy type, taxes per square foot, and state and loan-origination year fixed effects. The left panel shows the full sample. The right panel splits by borrower type: Undiversified Datacenter owners are firms whose primary business is datacenter leasing, as defined in Section 3.

Table 1. Summary Statistics

Variable	N	Median	SD	Max
Panel A: Property characteristics				
Rentable building area (SF)	550	135,617	383,398	6,500,000
Rent per square foot (USD)	507	13.1	11.6	201.4
Taxes per square foot (USD)	550	3.0	5.4	89.1
Years since origination	550	3.0	7.1	33.0
Panel B: Leverage outcomes				
Debt per square foot (USD)	550	358	5,003	62,414
Loan-to-value ratio	550	0.6	26.4	368.6
Panel C: Debt per SF by owner type (USD)				
Undiversified Datacenter owners	134	780		
Diversified owners	416	287		
Panel D: Debt per SF by bank OTD type (USD)				
High-OTD banks	264	876		
Low-OTD banks	183	142		

Note: This table reports summary statistics for datacenter properties with mortgage records. Panels A and B report medians, standard deviations, and maxima. Panels C and D report medians only. Debt per square foot is mortgage balance divided by rentable building area. Rent per square foot refers to base shell rent. Loan-to-value ratio is mortgage balance divided by property collateral value. Years since origination is the number of years between loan origination and 2025. Panel D excludes non-bank and government lenders.

Table 2. High-OTD Banks and Leverage

	ln(DebtPSF)	ln(DebtToRent)	ln(LTV)
High-OTD bank	1.223***	1.279***	1.217***
	(0.289)	(0.286)	(0.292)
Controls	Yes	Yes	Yes
State fixed effects	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes
R^2	0.372	0.360	0.362
Observations	550	550	550

Note: This table reports estimates from Equation (2). All specifications 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 3. High- vs. Low-OTD Banks

	ln(DebtPSF)	ln(DebtToRent)	ln(LTV)
High-OTD bank	1.605***	1.638***	1.586***
	(0.439)	(0.443)	(0.464)
Low-OTD bank	0.625	0.587	0.604
	(0.423)	(0.428)	(0.467)
Controls	Yes	Yes	Yes
State fixed effects	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes
R^2	0.378	0.365	0.366
Observations	550	550	550

Note: This table reports pooled regressions including both High-OTD and Low-OTD indicators. The omitted category is non-bank and government lenders. All specifications include the baseline controls and fixed effects described in Table 2. Standard errors clustered at the lender level are reported in parentheses. Tests of equality between the High-OTD and Low-OTD coefficients yield p-values of 0.0005, 0.0001, and 0.0007 for Columns (1)–(3), respectively. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 4. OTD and Borrower Selection

	DV = Undiversified Datacenter		
OTD score (z)	0.058**	0.068**	0.060**
	(0.025)	(0.027)	(0.027)
Controls	No	No	Yes
State fixed effects	No	Yes	Yes
Year fixed effects	No	Yes	Yes
R^2	0.378	0.342	0.332
Observations	326	326	326

Note: This table reports linear probability models relating lender OTD intensity to borrower type. The dependent variable equals one for properties owned by Undiversified Datacenter operators. Column (3) includes the baseline controls and fixed effects described in Table 2. Standard errors clustered at the lender level are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 5. High vs. Low-OTD: Matching Estimates

	ln(DebtPSF)		ln(LTV)	
	Prop. Score	Exact	Prop. Score	Exact
High vs. Low	0.761**	0.797**	0.817**	0.828**
	(0.319)	(0.353)	(0.357)	(0.375)
Controls	Yes	Yes	Yes	Yes
State fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes
Overlap	Yes	Yes	Yes	Yes
R^2	0.319	0.420	0.302	0.381
Observations	372	357	373	357

Note: This table compares High-OTD and Low-OTD bank loans under observable overlap. Propensity-score columns use inverse probability weighting. Exact-matching columns use coarsened exact matching based on state and property-size quartiles. Outcome regressions include the baseline controls and fixed effects described in Table 2. Standard errors clustered at the lender level are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 6. Excluding Large Banks

	ln(DebtPSF)	ln(DebtToRent)	ln(LTV)
High-OTD bank	1.457***	1.484***	1.375***
	(0.291)	(0.300)	(0.327)
Controls	Yes	Yes	Yes
State fixed effects	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes
R^2	0.443	0.420	0.415
Observations	381	381	381

Note: This table reports estimates from Equation (2) excluding the 16 banks with total assets exceeding $100 billion as of 2021 Q4, all of which are classified as High-OTD. All specifications include the baseline controls and fixed effects described in Table 2. Standard errors clustered at the lender level are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 7. Instrumental Variable Estimates

	OLS		2SLS: Balance Sheet IV		2SLS: Income IV	
	ln(DebtPSF)	ln(LTV)	ln(DebtPSF)	ln(LTV)	ln(DebtPSF)	ln(LTV)
High-OTD bank	0.771**	0.835**	1.258***	1.229***	1.288***	1.283***
	(0.324)	(0.330)	(0.408)	(0.436)	(0.397)	(0.428)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
State fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	326	326	326	326	326	326
First stage diagnostics						
Excluded instruments	—		Loan-to-assets ratio & Trading-assets ratio		Nonint. income share & Loan-to-assets ratio	
Kleibergen-Paap F	—		47.70		32.67	
Stock-Yogo (10%)	—		19.93		19.93	

Note: This table reports instrumental variable estimates. Columns (1)–(2) report OLS estimates restricted to the bank subsample with available 2021 Call Report data. Columns (3)–(4) use loan-to-assets and trading-assets-to-assets ratios as excluded instruments. Columns (5)–(6) use noninterest income share and loan-to-assets ratio as excluded instruments. All instruments are measured using 2021 Call Report data, before the 2022 rate-hike cycle and before the acceleration in datacenter lending. 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 November 2025. 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.3).

Lender originate-to-distribute (OTD) measures. We link lenders to regulatory Call Report data to obtain OTD-based distribution orientation measures. Section 3 describes construction of the OTD score and the High-OTD versus Low-OTD classification.

A.2 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 owner 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. 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 taxes per square foot using the mean of observed `taxespersf` among datacenters in the same county.

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.

A.3 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.

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 property collateral value. Collateral value is the maximum of three approaches: government tax assessment from matched parcel records, an income-based value capitalizing CoStar shell rent, 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

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

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 $2025 - \text{OriginationYear}_i$.

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.

A.4 Owner and Lender Classification

Owner types. We classify owners into four groups: Undiversified Datacenters (specialized datacenter landlords), Big Tech, Government, and Other private owners. The classification uses standardized owner names and owner group labels derived from CoStar ownership fields. The indicator `UndiversDC` equals one for properties owned by Undiversified Datacenter owners.

Lender mapping and lender types. The loan records report a lender name (originator). We standardize lender names and map them to lender identifiers used for the regulatory OTD measures. Lenders are assigned to High-OTD banks, Low-OTD banks, non-bank lenders, or government/cooperative lenders. When a lender cannot be matched reliably, the observation is excluded from specifications requiring OTD-based classification.

Originate-to-distribute (OTD) measures. The OTD score is constructed as described in Section 3 using Call Report data for 2022–2025. Banks in the top two deciles of average score are classified as High-OTD. Appendix Table B2 reports group means.

Bank size variables. We identify large banks as institutions with total assets exceeding $100 billion as of 2021 Q4, using Call Report data for domestic banks and publicly available balance sheet data for foreign institutions. This threshold captures 16 institutions, all classified as High-OTD, used in Table 6. We also construct $\ln(\text{Assets}_{2021})$ as a continuous predetermined size control used in Appendix Table C4.

Appendix B: Lender Classification and Validation

Table B1. Lender Classification

High-OTD lenders	**Low-OTD lenders**	**Non-bank lenders**	**Government lenders**
Toronto Dominion Bank	Truist Bank	New York Life Insurance	CoBank ACB
Goldman Sachs	U.S. Bank	Acore Capital Mortgage	Farm Credit Mid-America FLCA
Wells Fargo	Citizens Bank	American General Life Insurance	Farm Credit of the Virginias ACA
JPMorgan Chase	Huntington National Bank	Principal Life Insurance	Badgerland Financial FLCA
Sumitomo Mitsui Banking	Société Générale	AIG Asset Management US	Compeer Financial
Citibank	KeyBank	Fidelity Guaranty Life Insurance	Small Business Administration
Morgan Stanley	SunTrust Bank	Column Financial	
BNP Paribas	Hancock Whitney Bank	Argentic Real Estate Finance	
Deutsche Bank	First Citizens Bank & Trust	Thorofare Asset Based Lending	
Bank of America	Bank of Montreal	Guggenheim Real Estate	
Barclays	Mizuho Bank	Ladder Capital Finance	
Fifth Third Bank	First Citizens Bank	Midcap Financial Trust	
CIT Bank N.A.	Standard Chartered Bank	Sun Life Assurance Company	
Royal Bank of Scotland	Webster Bank	Nationwide Life Insurance	
MUFG Bank	Prosperity Bank	*Other non-bank lenders*	
UBS Real Estate Securities	Sunflower Bank		
Lehman Brothers Holdings	Pacific Western Bank		
BOK Financial	Eurohypo AG New York Branch		
Commerce Bank	Royal Bank of Canada		
	CIBC		
	Natixis New York Branch		
	BankUnited		
	Compass Bank		
	Blue Sky Bank		
	SouthTrust Bank		
	Valley National Bank		
	Texas Capital Bank		
	Simmons Bank		
	PNC Bank		
	Other Low-OTD lenders		

Note: This table reports the lender classification used in the analysis. Banks are grouped into High-OTD and Low-OTD categories using the Call Report balance-sheet score defined in Equation 1. Non-bank lenders include insurance companies, private credit funds, and other non-depository lenders. Government lenders include public and cooperative credit institutions.

Table B2. OTD Balance Sheet Measures

	High-OTD Banks	Low-OTD Banks	Difference (High−Low)
Loan assets/assets (%)	40.2	69.7	−29.6***
Repo/assets (%)	8.0	0.3	7.8***
Trading assets/assets (%)	4.9	0.1	4.8***
OTD score	6.8	−0.6	7.4***

Note: This table reports group means of regulatory ratios used to construct the bank-level OTD score over 2022–2025. Loan assets, trading assets, and repo assets are expressed as percentages of total assets. The OTD score combines standardized balance sheet components as defined in Equation 1. The final column reports High-Low differences in means. Stars indicate statistical significance from unequal-variance tests of equality. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table B3. DealScan Validation

	Lead Arranger	Term Loan B (TLB)	Full Syndication	Real Estate Loan
High-OTD bank	0.188***	0.057**	−0.013	−0.009
	(0.065)	(0.022)	(0.010)	(0.009)
Log deal size	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes
Observations	4,352	4,352	4,352	4,352

Note: This table reports linear probability models using DealScan lender-rows matched to the CoStar lender universe. Lead Arranger equals one if the bank acts as lead arranger; Term Loan B equals one if the deal includes an institutional term loan tranche; Full Syndication equals one if the loan is fully distributed; Real Estate Loan equals one if the deal is classified as a real estate loan. The omitted category is Low-OTD banks. All specifications include log deal size and loan-origination year fixed effects. Standard errors clustered at the lender level (32 clusters) are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Appendix C: Additional Results and Robustness

Table C1. Low-OTD Banks and Leverage

	ln(DebtPSF)	ln(DebtToRent)	ln(LTV)
Low-OTD bank	−0.470*	−0.530**	−0.478*
	(0.263)	(0.262)	(0.279)
Controls	Yes	Yes	Yes
State fixed effects	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes
R^2	0.338	0.324	0.332
Observations	550	550	550

Note: This table reports estimates from Equation (2), replacing the High-OTD indicator with a Low-OTD indicator. All specifications include the baseline controls and fixed effects described in Table 2. Standard errors clustered at the lender level are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table C2. Single-Property Loans

	ln(DebtPSF)	ln(DebtToRent)	ln(LTV)
High-OTD bank	1.055***	1.115***	1.052***
	(0.382)	(0.380)	(0.381)
Controls	Yes	Yes	Yes
State fixed effects	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes
R^2	0.368	0.356	0.365
Observations	379	379	379

Note: This table reports estimates from Equation (2) restricted to loans collateralized by exactly one property, eliminating balance allocation across collateral properties. All specifications include the baseline controls and fixed effects described in Table 2. Standard errors clustered at the lender level are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table C3. OTD Leverage Premium by Borrower Type

	ln(DebtPSF)	ln(DebtToRent)	ln(LTV)
OTD score	0.092	0.109	0.068
	(0.097)	(0.094)	(0.120)
OTD score × Undiversified DC	0.157*	0.165*	0.191*
	(0.094)	(0.094)	(0.111)
Implied OTD effect for Undiversified DC owners			
	0.249**	0.274***	0.259**
	(0.099)	(0.103)	(0.102)
Controls	Yes	Yes	Yes
State fixed effects	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes
R^2	0.456	0.448	0.420
Observations	326	326	326

Note: This table reports interaction regressions using the standardized continuous OTD balance-sheet score. OTD score × Undiversified DC is the product of the continuous OTD score and an indicator for Undiversified Datacenter owners. The implied OTD effect for Undiversified DC owners is the sum of the OTD score coefficient and the interaction coefficient, computed via `lincom`. All specifications include the baseline controls and fixed effects described in Table 2. Standard errors clustered at the lender level are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table C4. Predetermined Bank Assets

	ln(DebtPSF)	ln(LTV)
	(1)	(2)
High-OTD bank	0.760**	0.771**
	(0.378)	(0.377)
ln($Assets_{2021}$)	0.003	0.019
	(0.062)	(0.065)
Controls	Yes	Yes
State fixed effects	Yes	Yes
Year fixed effects	Yes	Yes
R^2	0.453	0.423
Observations	326	326

Note: This table reports estimates from Equation (2) augmenting the baseline specification with each bank's total balance-sheet assets in 2021 Q4, measured before the datacenter lending boom and the 2022 rate-hike cycle. All specifications include the baseline controls and fixed effects described in Table 2. Standard errors clustered at the lender level are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table C5. Oster (2019) Sensitivity

	ln(DebtPSF)	ln(DebtToRent)	ln(LTV)
$\hat{\beta}$ (controlled)	1.612	1.648	1.583
β^* ($\delta = 1$)	1.499	1.565	1.475
δ (s.t. $\beta = 0$)	4.22	4.55	4.54

Note: This table reports Oster (2019) sensitivity tests using $R_{\max} = 1.3 \times \hat{R}^2$. $\hat{\beta}$ is the controlled estimate, β^* is the bias-adjusted estimate at $\delta = 1$, and δ is the degree of selection on unobservables relative to observables required to reduce the coefficient to zero.

Appendix D: Instrumental Variable Validation

Table D1. IV Persistence and First Stage Strength

	2021 vs. 2022	2021 vs. 2023	2021 vs. 2024	2021 vs. 2025
Panel A: Instrument persistence (cross-bank correlations)				
Loan-to-assets ratio	0.941	0.921	0.901	0.867
Trading-assets ratio	0.995	0.989	0.975	0.971
Panel B: First stage (2022–2025 baseline, $N = 326$)				
Kleibergen-Paap F		47.70		
Stock-Yogo (10%) critical value		19.93		

Note: This table reports instrument persistence and first-stage strength for the baseline IV specification. Panel A reports cross-bank correlations between each instrument's 2021 value and its value in subsequent years, supporting the predetermination argument. Panel B reports the Kleibergen-Paap F-statistic from the first-stage regression using 2021 loan-to-assets and trading-assets-to-assets ratios as excluded instruments.

Table D2. Pre-Boom Placebo Test

	2015–2019 Pre-boom		2022–2025 Baseline	
	ln(DebtPSF)	ln(LTV)	ln(DebtPSF)	ln(LTV)
High-OTD bank	0.890	0.890	1.258***	1.229***
	(0.826)	(0.905)	(0.408)	(0.436)
Controls	Yes	Yes	Yes	Yes
State fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes
Observations	56	56	326	326
First stage diagnostics				
Kleibergen-Paap F	4.01	4.01	47.70	47.70
Stock-Yogo (10%)	19.93	19.93	19.93	19.93

Note: This table reports 2SLS placebo estimates applying the same 2021 instruments to datacenter loans originated before the lending boom. In the 2015–2019 pre-period the first stage is weak and the second-stage estimates are statistically insignificant, while the 2022–2025 baseline produces a strong first stage and significant estimates, consistent with the OTD channel becoming relevant only during the post-2021 lending boom. Standard errors clustered at the lender level are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table D3. Borrower Selection: OTD and Undiversified Datacenter Owners

	(1) Raw	(2) FE only	(3) FE + Controls	(4) 2SLS: BS IV	(5) 2SLS: Income IV
OTD score (z)	0.058**	0.068**	0.060**	0.051**	0.041
	(0.025)	(0.027)	(0.027)	(0.025)	(0.026)
Controls	No	No	Yes	Yes	Yes
State fixed effects	No	Yes	Yes	Yes	Yes
Year fixed effects	No	Yes	Yes	Yes	Yes
Observations	326	326	326	326	326
First stage diagnostics					
Excluded instruments		—		Loan-to-assets & Trading-assets	Nonint. income & Loan-to-assets
Kleibergen-Paap F		—		39.36	27.81

Note: This table reports OLS and 2SLS estimates of the effect of OTD propensity on the probability of financing an Undiversified Datacenter owner. OTD score (z) is the standardized continuous balance-sheet score. Columns (1)–(3) report OLS estimates without controls, with fixed effects, and with full controls. Columns (4)–(5) use 2021 Call Report balance-sheet measures as excluded instruments, before the 2022 rate-hike cycle and the acceleration in datacenter lending. Standard errors clustered at the lender level are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table D4. Funding Cost Diagnostics

	No brokered deposit control		Brokered deposit control	
	ln(DebtPSF)	ln(LTV)	ln(DebtPSF)	ln(LTV)
High-OTD bank	1.244***	1.122***	0.974*	1.011*
	(0.418)	(0.407)	(0.548)	(0.575)
Brokered deposit ratio			0.265	0.265
			(0.168)	(0.168)
Controls	Yes	Yes	Yes	Yes
State fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes
Observations	323	323	323	323
First stage diagnostics				
Kleibergen-Paap F	53.88	54.67	19.77	19.77
Stock-Yogo (10%)	19.93	19.93	19.93	19.93

Note: This table reports 2SLS estimates with and without a control for the brokered deposit ratio, restricted to the observations with non-missing brokered deposit data. The brokered deposit coefficient is statistically insignificant and the High-OTD estimates remain economically large, ruling out a funding cost interpretation. Standard errors clustered at the lender level are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table D5. Instrument Consistency

Instrument set	ln(DebtPSF)	ln(LTV)	KP F
Single instruments			
Loan-to-assets ratio only	1.452***	1.440***	30.13
	(0.423)	(0.451)	
Trading-assets ratio only	0.781	0.710	23.64
	(0.735)	(0.746)	
Noninterest income share only	2.714**	2.645**	6.73
	(1.066)	(1.066)	
Combined instruments			
Balance-sheet IVs	1.258***	1.229***	47.70
	(0.408)	(0.436)	
Income-based IVs	1.288***	1.283***	32.67
	(0.397)	(0.428)	
Controls	Yes	Yes	Yes
State fixed effects	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes
Observations	326	326	326
Overidentification			
Sargan-Hansen	Not computable—singleton fixed-effect cells		

Note: This table compares 2SLS estimates across alternative instrument sets. Single-instrument rows use each excluded instrument separately. Balance-sheet IVs use 2021 loan-to-assets and trading-assets-to-assets ratios. Income-based IVs use 2021 noninterest income share and loan-to-assets ratio. Convergence of just-identified estimates across instrument sets serves as a consistency check in place of a formal Sargan-Hansen statistic, which is not computable due to singleton fixed-effect cells. Standard errors clustered at the lender level are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.