Determinants of Pandemic-era CRE Distress: Implications for the Banking Sector*

David Glancy† Robert Kurtzman‡
Federal Reserve Board Federal Reserve Board

June 13, 2025

Abstract

We develop a two-step process to examine U.S. banks' exposure to pandemic-era commercial real estate (CRE) market strains and the extent to which unobserved behavior such as evergreening could be masking distress. We first use loan-level data on large banks' CRE portfolios to identify recent drivers of loan distress. We then examine the near-universe of CRE loan transactions to assess banks' exposure to high-default-risk CRE loans. Using several machine-learning algorithms, we demonstrate that cross-lender differences in CRE loan performance are broadly attributed to observable loan characteristics. In particular, small banks' strong performance reflects low holdings of at-risk office loans.

Keywords: commercial real estate, banks, CMBS

JEL Classification: G21, G23, R33

*We thank Vrinda Mittal (discussant), Joe Nichols, and seminar participants at the 2025 Day Ahead Conference on Financial Markets and Institutions, the Federal Reserve Board R&S workshop, the Philadelphia Fed, and the 2024 CREDA Real Estate Research Symposium for helpful comments. The views expressed in this paper are solely those of the authors and do not necessarily reflect the opinions of the Federal Reserve Board or anyone in the Federal Reserve System.

†Principal Economist, Division of Monetary Affairs, Federal Reserve Board, david.p.glancy@frb.gov

‡Principal Economist, Division of Research & Statistics, Federal Reserve Board, robert.j.kurtzman@frb.gov

1. INTRODUCTION

Higher interest rates and shifts in where people work and shop have created significant stress in pockets of the commercial real estate (CRE) market. Policy makers and academics have expressed concern that these developments could exacebate recent banking-sector strains, particularly for smaller banks, which frequently have significant CRE loan exposures (Acharya et al., 2023; Jiang et al., 2023; Board of Governors of the Federal Reserve System, 2023). While CRE loan delinquencies at small banks remained low through 2023, strains might be masked if small banks are engaged in evergreening (or colloquially, "extend and pretend") (Crosignani and Prazad, 2024; Cochrane and Seru, 2024; Peterson, 2024). Understanding whether small banks' loan performance reflects safer loans (as opposed to delayed loss recognition) is thus important for assessing the financial stability risks posed by CRE credit losses.

Analyzing such effects is complicated by a lack of detailed data on banks' CRE loan holdings. As a result of this data limitation, researchers have mostly relied on aggregate bank portfolio data (Faria-e Castro and Jordan-Wood, 2023, 2024), data from large banks (Crosignani and Prazad, 2024; Marsh and Pandolfo, 2024), or data from CRE segments with more public reporting (Gupta et al., 2022; Jiang et al., 2023) to assess risks posed to the banking sector. However, because banks serve a selected segment of the CRE market (Glancy et al., 2022a), extrapolating across different parts of the CRE market can be difficult. Indeed, Figure 1 demonstrates that small banks largely avoided the pandemic-era escalation in CRE delinquencies that plagued larger banks and securitized CRE loans (top panel). Moreover, the rise in delinquency at large banks is concentrated in office loans, further demonstrating the unevenness of CRE market strains (bottom panel). Understanding the implications of recent CRE stress requires a more detailed understanding of where loan performance is deteriorating and which lenders are exposed to troubled segments than can be found in the more widely available data.

In this paper, we compile data from a variety of sources to analyze why CRE loan performance differs across lenders, addressing this gap in the literature. The analysis proceeds in two steps.

First, we use loan-level panel data to investigate what factors are associated with CRE delinquency and why performance differs between bank and commercial mortgage-backed securities (CMBS) loans. To do so, we combine, harmonize, and analyze data from CMBS filings and confidential data from large banks' stress tests. We find that office loans held in CMBS pools are nearly 4 percentage points more likely to go delinquent than those held by banks. This effect is driven entirely by bank loans being secured by smaller properties, having lower loan-to-value (LTV) ratios, and being more likely to have recourse. For other property types, CMBS also underperform bank loans, but the difference is smaller and only partly attributable to those observable characteristics.[1]

In the second part of the paper, we investigate the role such observed characteristics play in driving the strong performance at small banks. Studying loan performance at large banks (i.e., banks with more than $100 billion in assets) is useful because (1) that is where performance has deteriorated more and (2) data availability enables more-detailed analysis. However, it is smaller banks that are highly exposed to CRE loans, as we show in Figure 2, which plots scatter plots of CRE-to-assets (left) and delinquency rates (right) by bank assets. Banks with between $1 billion and $10 billion in assets have the highest concentrations in CRE, but it is mostly banks with more than $100 billion in assets experiencing high delinquency rates. In effect, if troubles for CRE loans at CMBS and large banks portend similar problems at smaller banks, more significant stress may be looming. To understand such risks, we need to evaluate why CRE loan performance at small banks has remained comparatively resilient.

To address this question, we combine our analysis on the determinants of delinquency at large banks with transaction-level data covering a broader swath of the CRE market. While detailed loan-level panel data are generally unavailable for small banks, characteristics pertaining to loan sizes, property types, and locations are available due to the public reporting of mortgage liens. We can therefore estimate delinquency models with the large bank data and investigate the degree to which these observable factors can explain the strong performance at small banks. We do so using

[1] We also find that banks provide more extensions, which may contribute to these unobserved differences.

both easily interpreted models (OLS and low-complexity trees) where it is clear what variables drive performance differences and more flexible ones (K-nearest neighbors and random forests) that may better capture complex patterns in the data.

Regardless of the model chosen, we find that differences in CRE loan performance across bank sizes are mostly attributable to variation in loan size and the types and locations of the properties securing the loans. For the models that allow us to meaningfully decompose why performance differs, we find that, by far, the biggest difference is that small banks are less exposed to large-sized office loans, which performed particularly poorly in 2023. This factor alone accounts for about half of the 2 percentage point difference in delinquency rates between large and small banks.

Overall, these findings demonstrate that differences in loan performance across lenders are broadly attributable to observable characteristics of the loans they hold, rather than unobserved behavior such as evergreening. Large banks outperform CMBS due to differences in property sizes and underwriting terms, while small banks outperform large ones due to their low exposure to at-risk office loans. An important implication of these results is that it leaves little room for evergreening (or other ways banks might obfuscate distress) in explaining small banks' strong CRE performance. Conceptually, banks might outperform CMBS because capital preservation incentives cause them to evergreen (Crosignani and Prazad, 2024)—an action prohibited for CMBS.[2] Smaller banks in particular might be more likely to obscure losses due to lower examiner scrutiny (Hirtle et al., 2020) or a limited capacity to absorb losses (Peek and Rosengren, 2005). While these mechanisms could theoretically explain observed cross-lender differences in CRE loan performance and are certainly present to an extent, we show that first-order differences in performance are instead attributable to observable variation in loan underwriting and portfolio composition. These findings imply that the primary risk to small banks appears to come less from the CRE strains experienced through 2023 and more from the risk of stress spreading to other parts of the CRE market.

[2] While banks might allow an extension that causes higher eventual losses for the sake of delaying loss recognition, CMBS special servicers are obligated to maximize the net present value of the proceeds to the CMBS trust. Deviation from this servicing standard could subject them to litigation or ratings downgrades.

Our paper contributes to a growing body of work exploring the implications of pandemic-era shifts in where people live and work. Amid the shift to work from home, office owners have struggled to fill vacant or expiring space (Glancy and Wang, 2023), causing leasing revenue and office valuations to fall dramatically (Gupta et al., 2022). Central business districts (CBDs) have experienced particularly notable declines in commuting activity (Monte et al., 2023), commercial rents (Rosenthal et al., 2022), and office valuations (Ghosh et al., 2022). We document how these strains have passed through to CRE loan performance. While these factors have contributed to delinquency, the effects are mostly confined to a small segment of the market: large office loans, particularly in high work-from-home markets.

We also contribute to the literature analyzing banking-sector strains following Silicon Valley Bank's collapse. Most closely related is Jiang et al. (2023), who argues that potential CRE credit losses put many (mostly smaller) banks at risk of insolvency. This risk is compounded by the other strains stemming from interest rate increases and deposit runs in the spring of 2023 (Jiang et al., 2024; Cipriani et al., 2024; Choi et al., 2023). Crosignani and Prazad (2024) present evidence that extensions of troubled loans at large banks crowd out new lending. We show that while small banks are highly exposed to CRE loans, these risks are mitigated by the fact that their loans that are generally performing better in the current environment.

Finally, we contribute to the literature discussing banks' vulnerability to CRE stress more generally. This work demonstrates that CRE losses frequently are a driving force behind banking crises. This fact has been established in the context of the Savings and Loan Crisis (Browne and Case, 1992; Fenn and Cole, 2008) and the Great Recession (Cole and White, 2012; Friend et al., 2013) as well as internationally (Herring and Wachter, 1999; Gan, 2007). However, the monitoring of CRE risks is hampered by a paucity of data relative to other markets, particularly for loans from small banks (CRE Data Alliance, 2007). Our methodology allows us to overcome these limitations and analyze small banks' exposure to at-risk CRE loans.[3]

[3]Our work is also related to Acharya et al. (2024), who similarly bring new data to bear regarding banks' exposure to CRE risks. They demonstrates that banks are exposed to CRE through loans to REITs in addition to direct CRE

The rest of the paper proceeds as follows. Section 2 identifies what factors are associated with delinquency at large banks and CMBS. Section 3 estimates a model of loan performance using only loan characteristics that are observable for small banks' CRE loans and examines the extent to which those factors can account for small banks' lower delinquency rates. Section 4 concludes.

2. CRE LOAN PERFORMANCE AT LARGE BANKS AND IN CMBS

This section uses loan-level bank and CMBS data to investigate the loan and property characteristics affecting CRE loan performance in 2023.

2.1. Data and Methodology

The data on bank loans come from FR Y-14Q filings (the data underlying bank stress tests), which provide loan-quarter information on loans with committed balances over $1 million from banks with more than $100 billion in assets. The CMBS data come from Morningstar, which compiles loan-month data from CMBS disclosures.[4] We classify lenders by who holds (rather than who originates) loans, so bank-originated loans in CMBS pools are considered CMBS loans.

Banks serve some CRE segments that CMBS do not (Ghent et al., 2019; Glancy et al., 2022a). To focus on areas where both lenders overlap, we restrict the bank sample to first lien, non-owner-occupied, non-construction CRE loans. Bank data are reported as of quarter-end, so the analysis in this section uses the sample of loans that were outstanding as of the end of 2022 (in order to look at the effects of upcoming loan maturities), whereas the analysis in Section 3 uses the sample of loans that were still on the balance sheet as of the end of 2023 (to align with the reporting from the Call Reports). More information on how we clean and harmonize these data are available in Appendix A.1.1.

loans.

[4]We exclude agency deals from the CMBS data as well as defeased or real estate owned loans. For bank loans, we only include first-lien loans against already-constructed properties to better align the sample with that of the CMBS market.

Though smaller banks tend to have higher CRE concentrations, the Y-14 sample has the advantage of covering the banks for which CRE loan performance has materially deteriorated. These data are thus useful for examining the observed factors causing CRE loans to go delinquent. Additionally, the detailed reporting of loan and property characteristics for this sample allows us to examine differences in loan performance for otherwise similar bank and non-bank CRE loans. Consequently, the analysis is also informative about the extent to which harder to quantify factors, such as servicing differences across lenders, contribute to loan performance.

To investigate how observable characteristics relate to loan performance, we estimate linear regressions of the form:

$$100 \times \text{Delinquent}_{i,23} = \beta_1 \text{CMBS}_i + \beta_2 \text{Maturing}_{i,23} + \beta_3 \text{Office}_i + \gamma' X_{i,23} + \varepsilon_i, \tag{1}$$

for the sample of CRE loans that were outstanding as of year-end 2022. $\text{Delinquent}_{i,23}$ is an indicator for whether loan i is delinquent as of the last 2023 observation, defined here as being past due, performing beyond its maturity date, or liquidated.[5] The main independent variables are indicators for whether loan i was in a CMBS pool, was scheduled to mature in 2023, and was secured by an office property. The coefficient on CMBS_i reflects the difference in delinquency rates for CMBS loans compared with those from large domestic banks.

$X_{i,23}$ is a set of controls that includes other property-type dummies (multifamily is the omitted category). In some specifications, we layer in additional controls to assess whether they can account for some of the difference in delinquency rates between large banks and CMBS. These additional controls include LTV, property size, a recourse indicator, and geographic characteristics (a CBD indicator and the share of jobs in a city that can be done remotely). In the most expansive specifications, we also add occupancy and income controls to account for unobserved factors relating to properties' financial performance.[6]

[5] The quarter of observations is either 2023:Q4 if the loan is active at year-end or the quarter the loan was paid off or liquidated otherwise.

[6] Income is measured by the debt yield (the ratio of net operating income to the outstanding loan balance), which

Observing how β_1 changes as extra controls are added provides information on how much these characteristics can account for differences in loan performance. Recent work on cross-lender differences in CRE underwriting and performance points to a few potentially relevant variables:

1. **Size and location**: CMBS specialize in larger (Ghent et al., 2019) and more urban (Glancy and Wang, 2023) loans, which may be disproportionately affected by recent strains due to tenant mix and geographic differences in remote-work patterns (Marsh and Pandolfo, 2024). These effects are reflected in the controls for the (log) at-origination property value and geographic characteristics.

2. **Skin-in-the-game**: CMBS loans have higher LTVs (Glancy et al., 2022a) and are non-recourse (Glancy et al., 2023), so it may take smaller property value declines to induce default. Such effects are reflected in the at-origination LTV and recourse controls.

3. **Unobserved risks**: Information asymmetries associated with securitization can reduce the quality of CMBS loans (Ashcraft et al., 2019; Griffin and Priest, 2023). In the other direction, banks specialize in riskier loans because of their greater flexibility in managing distress (Black et al., 2017, 2020). The occupancy and income controls are meant to account for such differences in the risk of financial performance deteriorating.

4. **Servicing differences**: Banks can more easily modify distressed CRE loans (Black et al., 2017, 2020), which contributed to their better performance early in the pandemic (Glancy et al., 2022b). Banks may also extend troubled loans to preserve capital (Cipriani et al., 2024). In contrast, CMBS servicers are obligated to maximize recovery on a net present value basis and thus should only delay resolution if doing so increases the amount recovered. In a robustness exercise, we indeed show that banks indeed provide more extensions, though there is no clear reduction in maturity defaults as a result.

Summary statistics of the main variables of interest are shown in Table 1, with data for large bank

reflects the ability of a property's cash flows to pay off the loan. As income and occupancy would not get updated in the event that a loan pays off in 2023, we measure these financial variables as of a year before.

and CMBS loans shown in columns (1) and (2), respectively. As previously discussed, banks, on average, make smaller, lower LTV loans that often have recourse. The CMBS sample has an average delinquency rate about 3 percentage points higher than that for large banks. Much of this difference is due to CMBS liquidating more loans in 2023 and having more loans that are performing after their maturity date; the difference is only 1 percentage point when using a narrower definition that only counts past-due or nonaccrual loans as delinquent (which is more aligned with nonperforming loan (NPL) measures from the Call Reports). Columns (3) and (4) present the same data weighted by loan size so that the averages reflect aggregate portfolio shares. Though the patterns are similar to the unweighted results, delinquency rates are higher, reflecting worse performance of large loans.

2.2. Results

We present the main regression results in Table 2. Column (1) shows the most parsimonious specification, which only includes the CMBS, maturing loan, and property-type indicators. CMBS loans are about 1.7 percentage points more likely to become delinquent, office loans are about 3.3 percentage points more likely to become delinquent (relative to multifamily loans), and maturing loans are about 12.2 percentage points more likely to become delinquent. This last effect highlights the significant difficulty borrowers face in refinancing to pay off balloon loans; borrowers who are able to remain current over the life of the loan are frequently failing to pay off the loan as it comes due.

Column (2) adds in controls for property size and underwriting terms. Characteristics associated with bank loans—lower leverage, smaller properties, and recourse—are also associated with stronger loan performance. Adding these three additional controls reduces the coefficient on CMBS from 1.65 to 0.45, indicating that CMBS' tendency to make larger, higher-LTV, nonrecourse loans accounts for most of their inferior performance.

Column (3) adds in geographic controls for whether the property is in a CBD (according to CBRE market definitions) and the share of jobs in the MSA that can be done at home (as calculated by Dingel and Neiman, 2020). Column (4) adds in two variables pertaining to the property's financial situation: its occupancy rate and an indicator for whether the debt yield (net income as a share of the loan balance) is less than 8 percent. The directions of all of the effects are as would be expected, with CBD locations, telework exposure, and weaker financials all associated with worse loan performance. However, the additional controls do not dramatically change $\hat{\beta}_1$, which rises to about 0.9 when the financial controls are added.

Columns (5) to (8) repeat the same analysis but only for office properties. CMBS' underperformance is even more pronounced for office loans, with CMBS loans having a delinquency rate that is about 3.4 percentage points higher than that for bank loans. Additionally, defaults at maturity are even more prevalent for office loans, with a 2023 maturity raising the delinquency probability by nearly 20 percentage points.

Column (6) shows that the underperformance of CMBS for this segment can be entirely attributed to differences in loan sizes, LTVs, and the use of recourse. Once we control for these factors, CMBS office loans have delinquency rates that are slightly *lower* than those of bank loans. Put differently, we find that CMBS office loans perform somewhat better than non-recourse bank loans of similar loan sizes and LTVs, but they have higher delinquency rates overall due to these observable characteristics. While geographic risk factors (included in column 7) and performance metrics (included in column 8) affect the performance of office loans much more strongly than that of other loans, the inclusion of these controls does not change $\hat{\beta}_1$ much. These controls do, however, reduce the coefficient on $\ln(\text{Value at Orig.})$, indicating that high delinquency rates for large office properties partly reflect worse financial performance in high-telework markets.

What role do loan modifications play in these performance differences? One explanation for banks' modestly stronger performance relative to observably-similar CMBS loans is that banks are more willing or able to renegotiate stressed loans. For example, banks may prevent stressed borrowers from missing maturity payments by liberally providing extensions. A second, and equally important, question is *why* banks might modify more loans. If differences reflect superior resolution technology (e.g., fewer contractual frictions relative to modifying securitized loans), it would indicate that bank loans are actually safer. However, if differences reflect "extend and pretend" behavior—delaying loss recognition even if it means larger future losses—banks' lower delinquency rates would be little cause for optimism regarding eventual credit losses.

Table 3 provides evidence that extend and pretend behavior is not behind banks lower delinquency rates. Each specification presents regressions predicting delinquency (columns 1 and 2) or extension (columns 3—5) for the sample of loans that were slated to mature in 2023. Column (1) shows that about 16.5 percent of bank loans default at maturity, a rate that is 4 percentage points *higher* than for CMBS loans. Column (2) adds in the controls used in Table 2 and the difference between banks and CMBS gets slightly stronger (consistent with previous findings that bank loans have observable characteristics associated with stronger loan performance). The factors most predictive of maturity defaults are having office collateral and low debt yields, which are associated with default rates that are 9 and 8 percentage points higher, respectively. In short, banks' stronger loan performance does not reflect a lower rate of maturity default (i.e., the type of default that would be prevented by loan extensions).

While bank loans do not appear to be protected from maturity defaults, they do disproportionately receive extensions. Columns (3) and (4) predict whether loans received maturity date extensions in 2023 using the same specification. They show that 55 percent of maturing bank loans received extensions, compared to under 10 percent of CMBS loans, and these differences are not attributable to differences in other observable characteristics.

Why do banks provide more extensions? If banks extend loans because they have fewer impediments to renegotiation, this would induce them to modify less-troubled loans at the margin (Glancy et al., 2022b). If banks extend loans to delay loss recognition, they would disproportionately modify the lowest-quality loans, which would produce the highest losses absent intervention (Cipriani et al., 2024). To investigate which types of loans different lenders extend, column (5) adds an interaction between the CMBS dummy and dummy variables for two of the best predictors of maturity defaults: the low debt yield and office indicators. While banks are always predicted to provide more extensions, the difference is smaller for more-stressed loans.

In short, bank loans are more likely to be extended, but those marginal extensions are for less-stressed loans and they do not result in banks having fewer maturity defaults. These patterns could be explained by banks having lower frictions to modifying loans, causing borrowers likely to need extensions to sort into banks (Black et al., 2020) or causing borrowers to solicit opportunistic modifications that are not needed to prevent default (Glancy et al., 2022b). However, these patterns are generally inconsistent with the idea that banks are modifying their worst loans to hide default. This result is important for the next section; it indicates that large banks' delinquency rates are not contaminated by extend and pretend behavior, so low delinquency rates for a particular loan segment is more reasonably attributable to there being less stress in that segment.

To recap the findings, we determine that many of the drivers of CRE delinquency are consistent with our predictions given the nature of the stress. The largest driver is loan maturity, reflecting difficulty meeting balloon payments amid tighter credit conditions, lower valuations, and higher interest rates. Additionally, delinquency is higher for properties where demand for space has presumably weakened—namely, larger offices in telework-exposed markets.

Finally, the analysis points to a few proximate causes of delinquency that are potentially important for understanding the credit outlook for banks. First, large banks' CRE loans modestly outperform CMBS loans with similar terms, financial performance, and geographic characteristics. This result suggests that unobserved characteristics of banks loans—e.g., modification ability or borrowers' concerns about damaging existing bank relationships—support loan performance. Second, loans securing large properties have higher delinquency rates, suggesting that they have risk characteristics—e.g., a worse income outlook, difficulty making up operating shortfalls, or sponsors more willing to strategically default—that weigh on loan performance. The next section explores the extent to which such differences in observable loan characteristics can account for small banks' strong loan performance.

3. IMPLICATIONS FOR SMALL BANKS

To the extent that small banks lend against smaller and less urban properties than large banks, the patterns documented in Section 2 may contribute to their comparatively strong loan performance. To test this hypothesis, we compile data from various sources on the composition of CRE loan portfolios across different types of lenders, and estimate a set of models of bank loan performance using only the characteristics that are observable for small banks' CRE loan portfolios. We then examine the fitted delinquency rates to assess how much those factors can account for small banks' lower delinquency rates.

3.1. Composition of CRE Portfolios

The first step in understanding risk factors for small banks' CRE portfolios is to compile data on the composition of their CRE loan holdings. The primary data source we use is MSCI Real Capital Analytics (RCA), which sources data from both public records and industry contracts to provide detailed information on CRE transactions.

The main drawback to the RCA data is that it only covers properties above $2.5 million in value. Though this sample covers the majority of non-owner-occupied CRE lending, the omission of smaller transactions is potentially problematic given the strong association between property size and loan performance. To mitigate this potential bias, we supplement the RCA data with data on open commercial mortgage liens in the public records provided by CoreLogic. Specifically, we use RCA data for loans with original balances over $2.5 million (as such transactions should be reliably included in RCA) and CoreLogic data for loans below this threshold.[7]

To maintain consistency with the previous analysis, we continue to focus on non-owner-occupied CRE loans secured by existing properties—i.e., we exclude construction loans and owner-occupied

[7] Note that RCA's threshold is $2.5 million for *transactions*. We use loan balances and not transaction value because the Corelogic sample does not include the transaction value.

CRE loans. This sample covers the CRE loans that have exhibited the most stress through 2023, as can be seen in Figure 3. More information on how we construct the data, including details on lender name matching, sample selection, and how we impute whether RCA loans are still outstanding, is available in the Appendix A.1.2.

After harmonizing, name matching, and combining these data sources, we have a cross-section of at-origination characteristics (including lenders) for what should be the near-universe of outstanding commercial mortgages. For each of these loans i from a lender j (small bank, large bank, CMBS, and unclassified), we have a set of observable loan or property characteristics $X_{i,j}$ that includes the loan size, origination date, property type, and property location. While this set of variables does not include all of the factors studied in Section 2, it includes most of the major variables affecting performance that are likely to differ notably across lenders. The most significant change to the specification is adding loan size, $\ln(\text{Balance at Orig.})$, in place of the measures of LTV and property value. Because we do not reliably have information on property values for refinances in CoreLogic data (appraised values are available in Y-14Q and RCA data but not in public records), we replace the measures of LTV and property value with this single variable, which reflects both leverage and property size and is more reliably available in the data.

The weighted-average portfolio characteristics for large and small banks in these data are shown in the last two columns of Table 1. The averages for large banks are broadly in line with the weighted averages of the Y-14 data (shown in column 3), validating that RCA and CoreLogic data successfully match banks' true portfolio composition. The statistics show that loans at large banks tend to be larger in size, slightly more likely to be secured by office properties, and much more likely to be located in a CBD than loans at small banks.

3.2. *Bank Delinquency Models*

The second step in understanding differences in loan performance is to estimate the probability of delinquency as a function of the loan characteristics that are observable for the small bank sample. We use the Y-14 data to estimate a delinquency function $\hat{D}_m(X_{i,j}) = \mathbb{E}_m(\text{Delinquent}_{i,j}|X_{i,j}, j = \text{Large Bank})$ where m indexes the model used to estimate $D(\cdot)$. We only use bank data because we are most interested in applying the model to small banks' loans and thus want to minimize the extent to which unobservables—e.g., recourse usage or modification tendencies—differ across the estimation and prediction samples. We use a sample of loans and definition of delinquency that are aligned with Call Reports reporting to be able to benchmark the fitted predicted delinquency rates with actual ones. Namely, the sample considered is the pool of loans with outstanding balances as of 2023:Q4, and delinquency is measured by whether a loan is 30 or more days past due or nonaccrual.

In choosing a model to estimate the probability of delinquency, we face a tradeoff between interpretability and flexibility; simpler models are better for understanding why performance differs across lenders, whereas more complex ones may improve the fit and thus better assess the extent to which the variables considered can jointly account for differences in performance (with less clarity concerning which particular variables matter). To provide a mix of of these benefits, we estimate four models with an increasing degree of flexibility.

1. **Linear probability model (OLS)**: Specification includes property-type fixed effects, with the main other variables—CBD, teleworkable share, and ln(Balance at Orig.)—also interacted with the office indicator.

2. **Decision tree**: Sequentially splits the sample by feature values to achieve the best model improvement with each split. We set a high splitting threshold for an easily interpreted tree. The algorithm splits the feature space into a small number of regions and returns the Y-14 delinquency rate for the region (i.e., leaf) as a prediction.

3. **K-nearest neighbors (KNN)**: Finds the observations in the Y-14 data with the most similar characteristics, and the predicted delinquency rate is the average for the K-closest Y-14 loans, weighting by the similarity of the X vector.

4. **Random forest**: Takes bootstrapped samples of the Y-14 data and fits shallow trees to them. The predicted delinquency rate is the average prediction across the samples.

For the OLS estimator, the primary decision is the specification. Motivated by the results in Table 2, we choose a mostly linear model, with additional interactions of size, CBD and teleworkable share with the office indicator.

The latter three models estimate delinquency nonparametrically and account for interactions and nonlinearities without us specifying them. Instead, the primary decision is with regard to hyper-parameters. For each model, we search over a parameter grid and use stratified fivefold cross validation to choose the parameters that produce the lowest mean-squared error in the left out data. More details on these estimators and hyperparameter tuning are available in Appendix A.2.

The coefficient estimates for the linear model are shown in column (1) of Table 4. Variables pertaining to loan size and telework ability are demeaned so the coefficient on office shows the effect for non-CBDs when other risk factors are at their mean.

Overall, the results offer few surprises relative to what was found previously. Higher loan balances are associated with higher delinquency, particularly for offices, consistent with previous results showing higher delinquency rates for larger properties and high-LTV loans. Delinquency rates are also higher in CBDs and high-telework cities, particularly for offices.

The patterns are similar when defining delinquency using the broader definition from the previous analysis (column 2) or when predicting the year-ahead estimated probability of default based on banks' internal risk ratings (column 3). Thus, the loan characteristics associated with delinquency as reflected in the Call Reports are also associated with weaker performance more broadly considered, and the expectation of weaker performance going forward.

The decision tree has four nodes with the following predicted delinquency rates: non-office loans (0.5%), small office loans (1.4%), and large office loans in low- and high-telework-eligible markets (8.4% and 23%, respectively). Large office loans have at-origination balances greater than $23.3 million and high-telework-eligible markets have teleworkable share greater than 0.44.[8]

One potential source of bias to take into account in the tree in particular is that the way the model accounts for size effects is misspecified. The decision tree estimates only reflect whether office loans are above a particular size threshold, whereas Figure 4 indicates that the effects of size are continuous. The tree therefore is likely to underestimate delinquency for the very large loans (which get the higher weight in the portfolio aggregations), which may be why it produces estimates that are lower than the other models, which allow size effects to be more continuous.

The KNN and random forest estimators do not have as clear a correspondence between features and predicted delinquency. However, the fitted delinquency rates are highly correlated with the first two estimators, suggesting that common drivers are at play. Predictions from the tree estimator have correlations of 0.63 with the KNN predictions and 0.88 with the random forest predictions (see Appendix Figure A2). This result indicates the broad categorization in the tree estimates—where the primary division is between large office loans and everything else—drives much of the variation in the more complex estimates.

3.3. Cross-Bank Differences in Fitted Nonperforming Loans

What do these estimates imply for differences in loan performance across lenders? We now use the delinquency models discussed in Section 3.2 to generate lender-specific predicted delinquency rates using the portfolio data discussed in Section 3.1. The object of interest is $\text{Fitted Delinquency}_j = \sum_{i|j} \omega_{i,j} \hat{D}_m(X_{i,j})$, where $\omega_{i,j}$ is loan i's share in lender j's portfolio. These estimates show how much CRE portfolio compositions can account for cross-lender performance differences. If the drivers of CRE performance at small and large banks are similar but small banks just hold loans that are

[8]The decision tree is visualized in Appendix Figure A1.

safer on modeled dimensions, then fitted delinquency rates should match observed ones. If small banks' CRE loans perform better for other reasons (e.g., relationships, underwriting quality, or evergreening incentives), then their stronger performance would be for unobserved reasons and would not be reflected in Fitted Delinquency$_j$.

The top panel of Figure 5 plots observed delinquency rates by lender in the 2023:Q4 Call Reports (blue bars) and the fitted delinquency rates from the four delinquency models (the other bars). For loan types in the model, small and large banks have delinquency rates around 0.6 percent and 2.6 percent, respectively. This 2 percentage point differential is well explained by the loan and property characteristics included in models. Across the four models, the fitted delinquency rate for small banks ranges from 1.0 percent in the random forest model to 1.2 percent in the OLS model. In other words, about 1.4 to 1.6 of the 2 percentage point difference in delinquency—70 to 80 percent of the gap—can be attributed to differences in the CRE portfolio composition along just a couple of dimensions (loan size, location, and property type).

For large banks, the fitted delinquency rates align closely with the observed ones; they range from 2.19 percent to 2.63 percent, compared with an observed rate of 2.61 percent. Both the OLS- and random forest-fitted delinquency rates are within 2 basis points of the observed one. While this result is not particularly surprising given that estimates are fit to large bank data, it does increase confidence in the methodology. It would be possible for fitted delinquency rates to deviate from actual ones due to sampling problems with the RCA and CoreLogic data or residuals that are correlated with loan size (because the portfolio aggregations are weighted). That the predictions align with the observed data suggests that these are not major problems.To be more precise, our results would be biased if we over- or under-sampled loans in a way that correlated with loan performance (e.g., if smaller loans were under-sampled due to reporting issues). Table 1 indicates that this is not a problem, as portfolios at large banks in the RCA and CoreLogic data match those in the Y-14 data.

What drives these differences in fitted delinquency? The OLS- and tree-based predictions can be decomposed to clarify why loans at large and small banks appear to perform differently. The bottom panel of Figure 5 presents a waterfall chart showing the various factors contributing to differences in loan performance in the OLS model. The first red bar shows the residual, or the unexplained amount by which small banks overperform. Small banks have an OLS-fitted delinquency rate of 1.2%, meaning that about 0.6 percentage points of the overperformance is driven by unmodeled factors. The other bars show how much individual variables in the regression contribute to performance differences.[9] While small banks benefit from having smaller loans in general and fewer loans in troubled areas (CBDs or high-telework cities), by far the biggest component is the size-by-office interaction; this effect accounts for almost 1 percentage point of the 1.4 percentage point difference in fitted delinquency rates. Most of the other effects work in the same direction, but nothing else contributes more than 0.28 percentage points to the difference. Those other variables have either too small an effect on delinquency or too little difference across bank sizes to contribute as notably. In short, the OLS estimates indicate that half of small banks' superior CRE performance is due to their low exposure to large office loans.

A decomposition of the tree estimates tells a similar story. Table 5 provides the predicted delinquency rate for different loan segments (column 1) and the share of large and small banks' CRE portfolios in those segments (columns 2 and 3, respectively). While small banks have a modestly smaller office exposure (≈17.5%) relative to large banks (≈20%), the bigger differentiator is the size of the office loans. Roughly 15% of large banks' CRE loans are against large office loans (offices with an at-origination balance above $23.3 million) compared to only about 4% of small bank loans. As small office loans have a delinquency rate of 1.4% in the Y-14 data while large ones have delinquencies above 8% (and much above in high-telework areas), this composition produces large differences in the fitted delinquency rates.

[9] Specifically, if β is the vector of regression coefficients and $\overline{X}_j$ is the balance-weighted average vector of loan characteristics for lender type j, then the difference in OLS-fitted delinquency rates between lender j and j' is $\beta'(\overline{X}_j - \overline{X}_{j'})$. Each bar shows a particular $\beta_k(\overline{X}_{j,k} - \overline{X}_{j',k})$, where k indexes explanatory variables.

4. CONCLUSION

Rising interest rates and shifts in the demand for space have impaired the performance of many CRE properties. As banks are large holders of CRE loans, these developments have generated concern about CRE exposure exacerbating other banking-sector strains.

Using a combination of different sources, we shed light on the factors affecting loan performance across different types of lenders. CRE loans held by large banks were less likely to go delinquent in 2023 than those held by CMBS. These differences can mostly be accounted for by banks making smaller loans where borrowers have more skin in the game (from either recourse or property equity). Lower NPLs at small banks are estimated to predominantly reflect small banks' lower exposure to at-risk office loans (i.e., loans secured by larger office properties). This result indicates that small banks' comparatively strong performance mostly reflects differences in what loans they hold, rather than how they are serviced or how they are classified in accounting statements.

All told, these findings suggest that strains would need to expand to other CRE market segments to cause widespread distress at small banks. Of course, just because small banks' CRE loans have characteristics that have insulated them so far, they are not necessarily immune to future stresses. Some segments that have performed reasonably well could reasonably deteriorate in the future. Difficulties leasing new space or obtaining stabilized financing could create future problems for construction loans as more projects exit construction into a challenging environment. Additionally, multifamily delinquency is rising due to the effects of higher interest rates and operating expenses, along with competition from new supply.

Thus, while this study provides a thorough analysis of why small bank CRE loan performance has held up so far, and gives some reason for optimism about the outlook going forward, the situation warrants monitoring, especially if the CRE market starts showing signs of broader strains.

References

ACHARYA, V. V., M. GOPAL, M. JAGER, AND S. STEFFEN (2024): "Shadow Always Touches the Feet: Implications of Bank Credit Lines to Non-Bank Financial Intermediaries," *Available at SSRN 4847858*.

ACHARYA, V. V., A. SAVOV, AND P. SCHNABL (2023): "Chapter 3: Banks, Interest Rate Risk and Systemic Risk–Theoretical and Historical Perspectives," in *SVB and Beyond: The Banking Stress of 2023*, ed. by M. P. Richardson, K. L. Schoenholtz, and B. Tuckman, Volatility and Research Institute of New York University Stern School of Business, 35–72.

ASHCRAFT, A. B., K. GOORIAH, AND A. KERMANI (2019): "Does Skin-in-the-Game Affect Security Performance?" *Journal of Financial Economics*, 134, 333–354.

BLACK, L., J. KRAINER, AND J. NICHOLS (2017): "From Origination to Renegotiation: A Comparison of Portfolio and Securitized Commercial Real Estate Loans," *The Journal of Real Estate Finance and Economics*, 55, 1–31.

——— (2020): "Safe Collateral, Arm's-Length Credit: Evidence from the Commercial Real Estate Mortgage Market," *The Review of Financial Studies*, 33, 5173–211.

BOARD OF GOVERNORS OF THE FEDERAL RESERVE SYSTEM (2011-2024a): "Call Reports Data," .

——— (2011-2024b): "FR Y-14Q Schedule H.2, Capital Assessments and Stress Testing: Commercial Real Estate," .

——— (2023): "Financial Stability Report." .

BROWNE, L. E. AND K. E. CASE (1992): "How the commercial real estate boom undid the banks," in *Conference Series;[Proceedings]*, Federal Reserve Bank of Boston, vol. 36, 57–113.

CHOI, D. B., P. GOLDSMITH-PINKHAM, AND T. YORULMAZER (2023): "Contagion Effects of the Silicon Valley Bank Run," *NBER Working Paper No. 31772*.

CIPRIANI, M., T. M. EISENBACH, AND A. KOVNER (2024): "Tracing Bank Runs in Real Time," *Federal Reserve Bank of New York Staff Reports*, 1104.

COCHRANE, J. H. AND A. SERU (2024): "Ending Bailouts, at Last," *Journal of Law, Economics & Policy*, 19, 169—93.

COLE, R. A. AND L. J. WHITE (2012): "Déjà vu all over again: The causes of US commercial bank failures this time around," *Journal of financial services Research*, 42, 5–29.

CORELOGIC, INC (2000-2024): "CoreLogic Real Estate Data," .

CRE DATA ALLIANCE (2007): "Commercial Real Estate Data: Towards Parity with Other Asset Classes," *White paper*.

CROSIGNANI, M. AND S. PRAZAD (2024): "Extend-and-Pretend in the U.S. CRE Market," *Federal Reserve Bank of New York Staff Reports*, 1130.

DINGEL, J. I. AND B. NEIMAN (2020): "How Many Jobs can be Done at Home?" *Journal of Public Economics*, 189, 104235.

FARIA-E CASTRO, M. AND S. JORDAN-WOOD (2023): "'Stress Testing' Banks on Commercial Real Estate," *Economic Synopses*, 1–3.

——— (2024): "Commercial Real Estate Exposure and Bank Stock Returns," *St. Louis Fed, On the Economy*.

FENN, G. W. AND R. A. COLE (2008): "The Role of Commercial Real Estate Investments in the Banking Crisis of 1985-92," *Available at SSRN 1293473*.

FRIEND, K., H. GLENOS, AND J. B. NICHOLS (2013): “An analysis of the impact of the commercial real estate concentration guidance,” *Federal Reserve Board and Office of the Comptroller of the Currency*.

GAN, J. (2007): “The real effects of asset market bubbles: Loan-and firm-level evidence of a lending channel,” *The Review of Financial Studies*, 20, 1941–1973.

GHENT, A. C., W. N. TOROUS, AND R. I. VALKANOV (2019): “Commercial Real Estate as an Asset Class,” *Annual Review of Financial Economics*, 11, 153–71.

GHOSH, C., L. ROLHEISER, A. VAN DE MINNE, AND X. WANG (2022): “The Price of Work-from-Home: Commercial Real Estate in the City and the Suburbs,” *Working Paper*.

GLANCY, D., J. R. KRAINER, R. J. KURTZMAN, AND J. B. NICHOLS (2022a): “Intermediary Segmentation in the Commercial Real Estate Market,” *Journal of Money, Credit and Banking*, 54, 2029–80.

GLANCY, D., R. KURTZMAN, L. LOEWENSTEIN, AND J. NICHOLS (2023): “Recourse as Shadow Equity: Evidence from Commercial Real Estate Loans,” *Real Estate Economics*, 51, 1108–36.

GLANCY, D. AND J. C. WANG (2023): “Lease Expirations and CRE Property Performance,” *Federal Reserve Bank of Boston Research Department Working Papers No. 23-10*.

GLANCY, D. P., R. J. KURTZMAN, AND L. LOEWENSTEIN (2022b): “Loan Modifications and the Commercial Real Estate Market,” *FEDS Working Paper 2022-050*.

GRIFFIN, J. M. AND A. PRIEST (2023): “Is COVID Revealing a Virus in CMBS 2.0?” *The Journal of Finance*, 78, 2233–76.

GUPTA, A., V. MITTAL, AND S. VAN NIEUWERBURGH (2022): “Work from Home and the Office Real Estate Apocalypse,” *NBER Working Paper No. 30526*.

HERRING, R. J. AND S. M. WACHTER (1999): "Real estate booms and banking busts: An international perspective," *The Wharton School Research Paper*.

HIRTLE, B., A. KOVNER, AND M. PLOSSER (2020): "The Impact of Supervision on Bank Performance," *The Journal of Finance*, 75, 2765–808.

JIANG, E. X., G. MATVOS, T. PISKORSKI, AND A. SERU (2023): "Monetary Tightening, Commercial Real Estate Distress, and US Bank Fragility," *NBER Working Paper No. 31970*.

——— (2024): "Monetary Tightening and U.S. Bank Fragility in 2023: Mark-to-Market Losses and Uninsured Depositor Runs?" *Journal of Financial Economics*, 195, 103899.

MARSH, W. B. AND J. PANDOLFO (2024): "Banks' Commercial Real Estate Risks Are Uneven," *Federal Reserve Bank of Kansas City, Economic Bulletin*, 1–4.

MONTE, F., C. PORCHER, AND E. ROSSI-HANSBERG (2023): "Remote Work and City Structure," *American Economic Review*, 113, 939–81.

MORNINGSTAR, INC. (2011-2024): "Morningstar CMBS Data Product," .

MSCI (2001-2024): "Real Capital Analytics," .

PEEK, J. AND E. S. ROSENGREN (2005): "Unnatural Selection: Perverse Incentives and the Misallocation of Credit in Japan," *American Economic Review*, 95, 1144–66.

PETERSON, D. M. (2024): "U.S. Commercial Real Estate Is Headed Toward a Crisis," *Harvard Business Review*.

ROSENTHAL, S. S., W. C. STRANGE, AND J. A. URREGO (2022): "JUE Insight: Are City Centers Losing their Appeal? Commercial Real Estate, Urban Spatial Structure, and COVID-19," *Journal of Urban Economics*, 127, 103381.

Figure 1: Nonperforming Loan Rates over Time

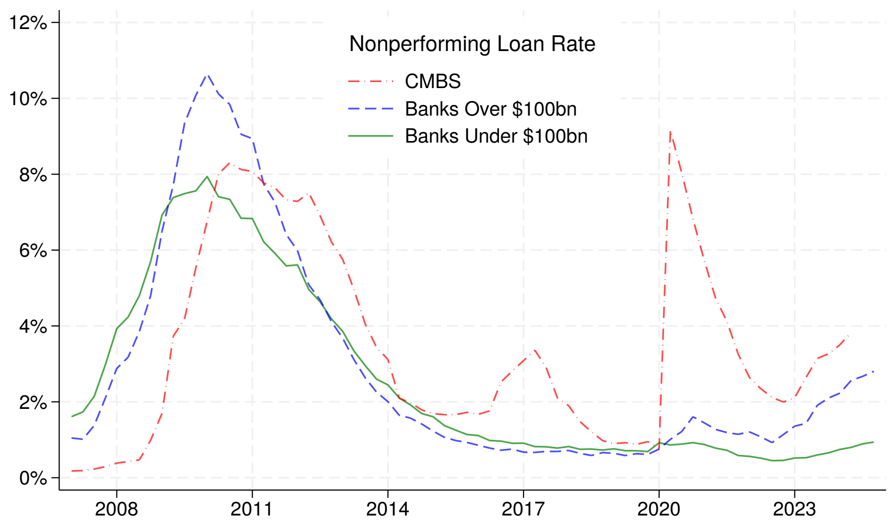

(a) Delinquency, by Lender

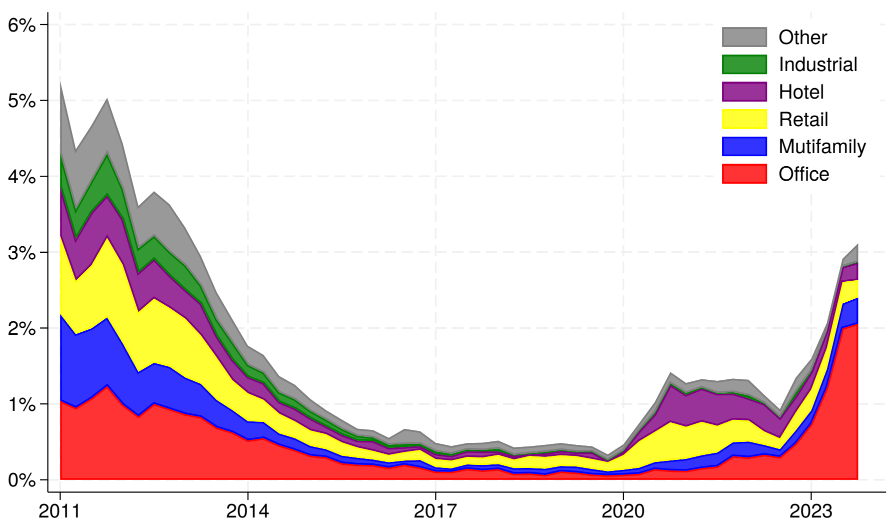

(b) Decomposition, by Property Type (Large Banks)

Notes: The top panel plots CRE nonperforming loan (NPL) rates over time for non-agency CMBS loans (red) and for CRE loans held by banks with more (blue) and less (green) than $100 billion in assets. NPL rates are loans that are 30 days or more past due or nonaccrual, plotted as a share of aggregate outstanding balances. The bottom panel uses Y-14 data to decompose the NPL rate for large banks' nonowner-occupied income-producing CRE loans by property type, with the shaded region showing the contribution of a particular property type.
Sources: Authors' calculations using the Call Reports and Morningstar (left) and the Y-14Q H.2 Schedule (right).

Figure 2: CRE Exposure, by Bank Size

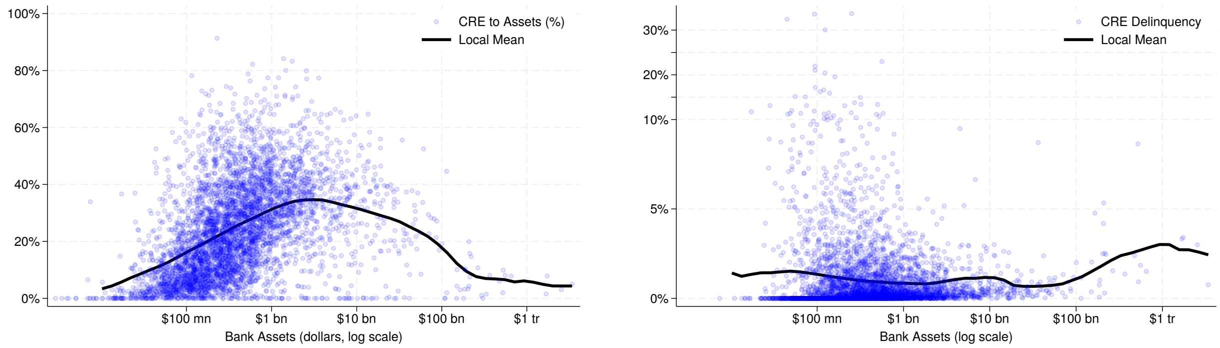

(a) CRE Concentrations

(b) CRE Delinquency

Notes: The figure plots CRE loans as a share of assets (left) and CRE delinquency rates (right) by bank size as of 2023:Q4. CRE includes non-farm non-residential, multifamily, and construction and land development loans. Blue dots report CRE shares or delinquency rates at individual banks, while the black line plots an estimate of these variables for banks of a given size (the kernel-weighted local mean). The scale for the y-axis changes in the right panel after 10% to improve the visibility of differences in performance of banks within typical ranges.
Sources: Authors' calculations using the Call Reports.

Figure 3: Nonperforming Loan Rates, by Bank Size

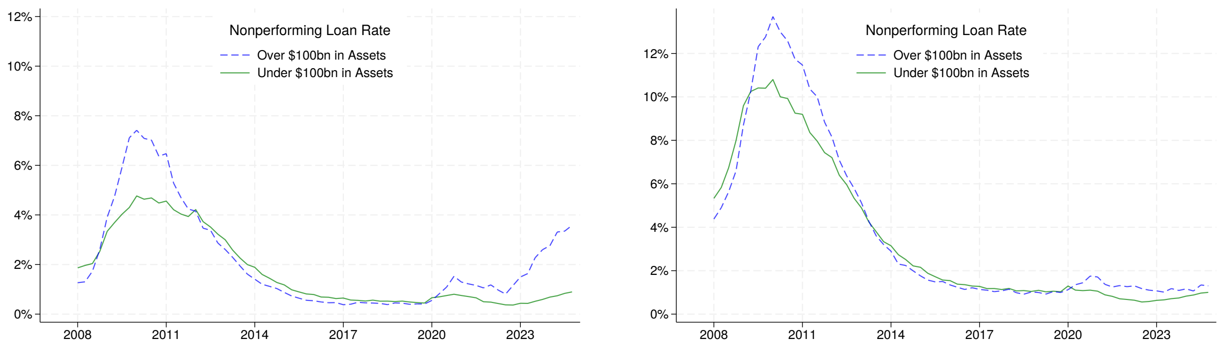

(a) Existing Investment Properties

(b) Other CRE Loans

Notes: The figure plots CRE nonperforming loans (NPLs) over time for banks above (blue) and below (green) $100 billion in assets. The left panel plots NPLs for the loan types included in the delinquency models (non-owner-occupied non-farm non-residential (NFNR) and multifamily loans), and the right panel plots NPLs for other CRE loans (owner-occupied NFNR and construction and land development loans). NPLs are loans that are 30 days or more past due or nonaccrual, plotted as a share of aggregate outstanding balances.

Sources: Authors' calculations using the Call Reports.

Figure 4: CRE Delinquency Rates, by Loan Size

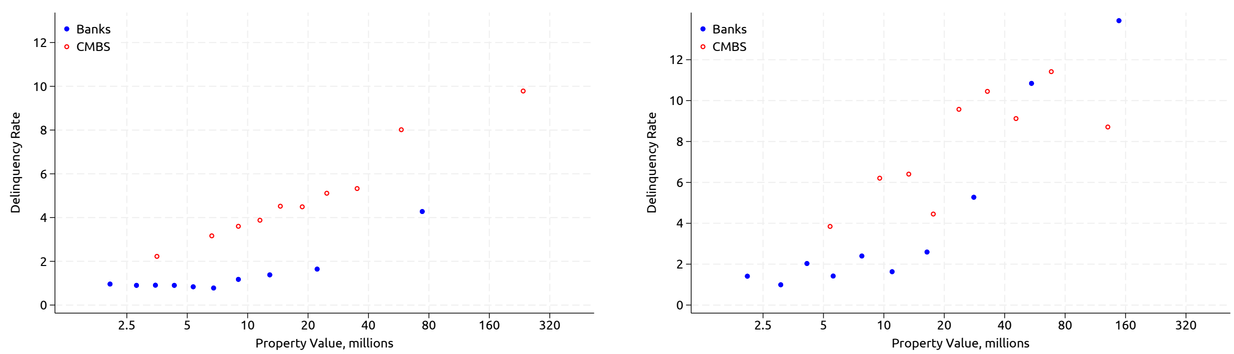

(a) Full Sample

(b) Office Loans

Notes: The figure plots binscatter estimates of 2023 CRE loan delinquency rates by the logarithm of the property value at origination. The red and blue dots are the delinquency rates for CMBS and Y-14 bank loans at different deciles of their respective property size distribution, respectively. The left panel presents delinquency rates pooling across property types, while the right panel restricts the sample to office properties.

Sources: Authors' calculations using the Y-14Q H.2 Schedule and Morningstar.

Figure 5: Realized versus Expected Nonperforming Loan Rates

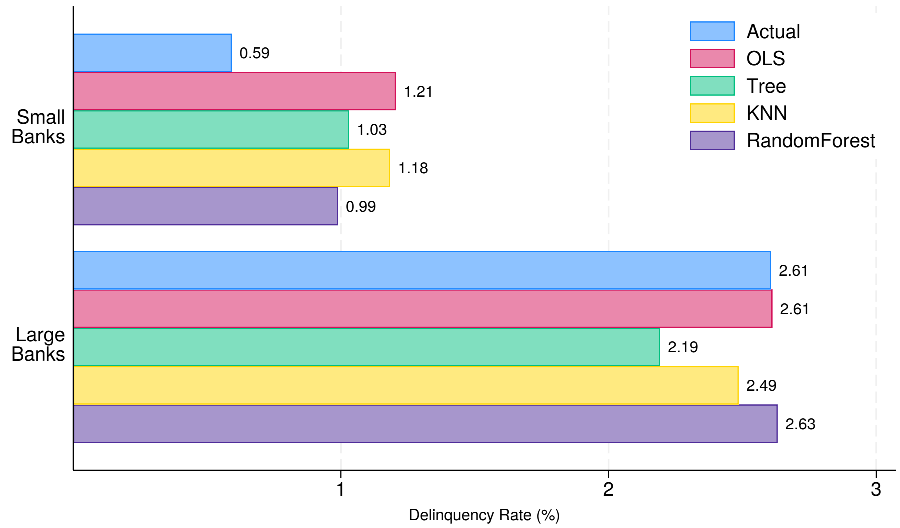

(a) Across Bank Sizes and Models

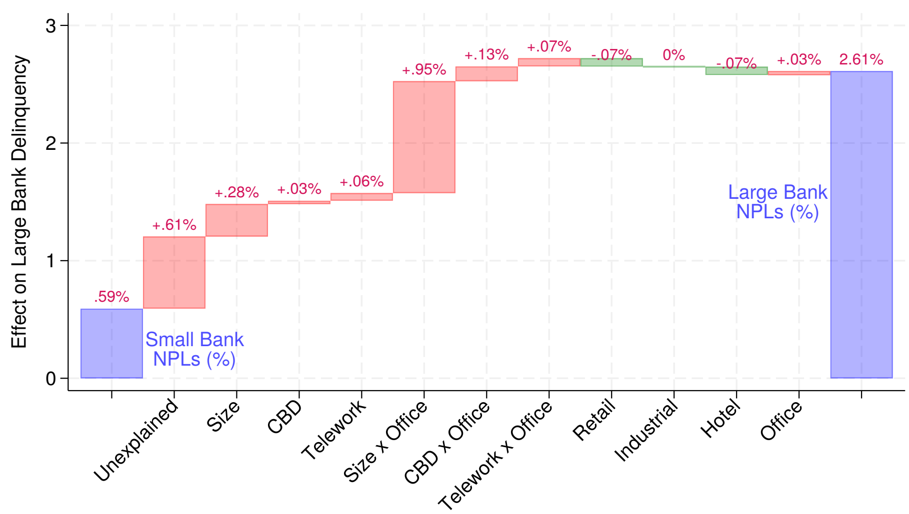

(b) OLS Decomposition of Difference Between Large and Small Bank NPLs

Notes: The top figure plots the the 2023:Q4 nonperforming loan rate (blue), and the weighted-average expected nonperforming loan rates based on a linear probability (red), a regression tree (green), a K-nearest neighbors (KNN)(yellow), and a random forest (purple) model. The top (bottom) set of bars pertains to banks with under (over) $100 billion in assets. The bottom figure decomposes differences in nonperforming loan rates between large and small banks. Blue bars show the rates for small (left) and large (right) banks. The bars in between show how much each variable from the regression in column (1) of Table 4 contributes to the difference (except for the first bar, which gives the unexplained component).
Sources: Authors' calculations using the Y-14Q H.2 Schedule, MSCI Real Capital Analytics, and CoreLogic.

Table 1: Average CRE Loan Characteristics across Samples

Data	Y-14	Morningstar	Y-14	Morningstar	RCA/CoreLogic	
Sample	Large Banks	CMBS	Large Banks	CMBS	Large Banks	Small Banks
	(1)	(2)	(3)	(4)	(5)	(6)
Delinquent$_{i,23}$	1.43	4.51	3.77	7.12		
Delinquent (Call defn.)	0.99	1.99	3.23	3.57	2.61	0.59
ln(Value at Orig.)	15.84	16.79	17.62	19.22		
LTV at Orig.	0.53	0.61	0.58	0.61		
Recourse	0.59	0.00	0.45	0.00		
Occupancy	0.94	0.89	0.90	0.87		
Debt Yield$<$.08	0.31	0.20	0.39	0.26		
ln(Balance at Orig.)	15.16	16.22	17.10	18.66	17.24	15.47
CBD	0.08	0.08	0.16	0.22	0.15	0.07
Teleworkable Share	0.39	0.37	0.39	0.37	0.39	0.37
Office	0.11	0.18	0.23	0.32	0.20	0.17
Retail	0.17	0.37	0.14	0.26	0.13	0.25
Industrial	0.08	0.06	0.12	0.10	0.17	0.20
Hotel	0.03	0.15	0.07	0.20	0.05	0.10
N	42267	15532	42267	15532	133954	209810
Weighted			✓	✓	✓	✓

Notes: Columns (1) and (2) present average loan characteristics for the sample of loans from large banks and CMBS that were outstanding as of the end of 2022 (the sample studied in Section 2). Columns (3) and (4) provide the same information, but weighting by the at-origination loan balance. Columns (5) and (6) present weighted-average loan characteristics for the sample of outstanding loans at large and small banks, respectively, based on the RCA and CoreLogic data. The measure of delinquency for the RCA and CoreLogic samples is the aggregate delinquency rate for nonowner-occupied nonfarm nonresidential and multifamily CRE loans from Call Reports since information on loan performance is generally unavailable in the original data.
Sources: Authors' calculations using the Y-14Q H.2 Schedule, Morningstar, CoreLogic, MSCI RCA, and the Call Reports.

Table 2: Loan Performance, by Lender Type

	$100 \times \text{Delinquent}_{i,23}$							
	Full Sample				Offices			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
CMBS	1.65**	0.45	0.69*	0.88**	3.36**	-2.17^{+}	-1.69	-1.01
	(0.26)	(0.36)	(0.33)	(0.33)	(0.67)	(1.19)	(1.18)	(1.16)
Maturing	12.23**	11.82**	11.77**	11.36**	19.56**	18.45**	18.16**	16.83**
	(0.96)	(0.95)	(0.95)	(0.95)	(1.88)	(1.86)	(1.85)	(1.80)
Office	3.37**	2.59**	2.48**	2.19**				
	(0.33)	(0.30)	(0.30)	(0.30)				
LTV at Orig.		4.87**	5.53**	3.92**		13.23**	15.22**	12.42**
		(0.72)	(0.83)	(0.81)		(1.88)	(1.98)	(1.89)
ln(Value at Orig.)		0.83**	0.73**	0.50**		1.97**	1.44**	1.04**
		(0.11)	(0.10)	(0.10)		(0.30)	(0.31)	(0.29)
Recourse		-0.40	-0.26	-0.30		-3.61**	-3.24**	-2.99**
		(0.25)	(0.23)	(0.23)		(1.06)	(1.05)	(1.02)
CBD			2.31**	1.87**			5.15**	3.92**
			(0.46)	(0.39)			(1.06)	(1.04)
Teleworkable Share			7.37**	5.24**			14.11**	11.77*
			(1.84)	(1.82)			(5.16)	(5.00)
Occupancy				-14.72**				-20.32**
				(1.61)				(2.53)
Debt Yield<.08				1.88**				5.79**
				(0.45)				(0.99)
R^2_a	0.056	0.061	0.064	0.081	0.072	0.093	0.100	0.134
Observations	57,799	57,799	57,799	56,873	7,652	7,652	7,652	7,505
Other Property Fixed Effects?	✓	✓	✓	✓				

Notes: This table presents estimates from the equation:

$$100 \times \text{Delinquent}_{i,23} = \beta_1 \text{CMBS}_i + \beta_2 \text{Maturing}_{i,23} + \beta_3 \text{Office}_i + \gamma' X_{i,23} + \varepsilon_i,$$

where $\text{Delinquent}_{i,23}$ is an indicator for whether loan i is delinquent as of the last observation in 2023 (2023:Q4 if the loan is active as of the end of the year, or the quarter the loan was paid off or liquidated otherwise). Loans that are liquidated or performing beyond their maturity date count as delinquent. The main independent variables of interest are whether loan i is in a CMBS pool, whether the loan was scheduled to mature in 2023, and whether the loan is secured by an office property. Fixed effects for other property types are included but not displayed (multifamily is the omitted category). Column (2) adds controls for whether the loan has recourse, the at-origination LTV, and the logarithm of the property value at origination. Column (3) adds controls for whether the property is in a CBD and the share of the city's employment that can be done at home (Dingel and Neiman, 2020). Column (4) adds controls for the occupancy and an indicator for whether the debt yield is less than 8% (both as of a year previously). Columns (5) to (8) repeat the same analysis but restrict the sample to office properties. Standard errors, in parentheses, are clustered by bank-origination year for bank loans and CMBS deal for CMBS loans. $^{+},^{*},^{**}$ indicate significance at 10%, 5%, and 1%, respectively.
Sources: Authors' calculations using the Y-14Q H.2 Schedule and Morningstar.

Table 3: Loan Modifications, by Lender Type

	$100 \times \text{Delinquent}_{i,23}$		$100 \times \text{Extension}_{i,23}$		
	(1)	(2)	(3)	(4)	(5)
CMBS	-4.07*	-5.71*	-46.45**	-43.15**	-45.33**
	(1.85)	(2.29)	(2.90)	(3.66)	(3.97)
...×Office					3.87
					(3.76)
...×Debt Yield<.08					5.90^{+}
					(3.41)
Office		9.48**		-1.68	-3.26
		(2.05)		(2.94)	(3.62)
ln(Value at Orig.)		0.79		10.21**	10.07**
		(0.56)		(0.69)	(0.69)
LTV at Orig.		19.13**		21.44**	21.75**
		(5.09)		(5.35)	(5.32)
Recourse		-3.59		15.02**	14.74**
		(2.33)		(3.43)	(3.43)
CBD		7.77**		2.68	2.50
		(2.41)		(2.45)	(2.46)
Teleworkable Share		20.10		-30.01*	-30.90*
		(12.62)		(13.60)	(13.67)
Debt Yield<.08		8.32**		7.10**	4.90^{+}
		(1.81)		(2.10)	(2.82)
Occupancy		-22.26**		11.91*	11.69*
		(5.00)		(5.47)	(5.48)
Constant	16.56**	-0.92	55.64**	-136.88**	-132.84**
	(1.42)	(11.95)	(2.62)	(13.89)	(13.94)
R^2_a	0.003	0.057	0.235	0.334	0.334
Observations	3,459	3,392	3,459	3,392	3,392
Other Property Fixed Effects?		✓		✓	✓

Notes: This table presents estimates from the equation:

$$100 \times \text{Extension}_{i,23} = \beta_1 \text{CMBS}_i + \beta_2 \text{Office}_i + \gamma' X_{i,23} + \varepsilon_i,$$

where $\text{Extension}_{i,23}$ is an indicator for whether loan i received an extension (i.e., had its maturity date pushed out) in the 2023. The sample includes loans that were outstanding as of the end of 2022 that were scheduled to mature in 2023. The main independent variable of interest are whether loan i is in a CMBS pool. Fixed effects for property types besides offices are included but not displayed (multifamily is the omitted category). Columns (1) and (2) predict whether loans go delinquent in 2023, while columns (3)–(5) predict whether they received maturity date extensions. Columns (2) and (4) add the controls from Table 2, while column (5) additionally includes interactions of the CMBS indicator with the office and low debt yield dummy variables. Standard errors, in parentheses, are clustered by bank-origination year for bank loans and CMBS deal for CMBS loans. $^{+}$,*,** indicate significance at 10%, 5%, and 1%, respectively.
Sources: Authors' calculations using the Y-14Q H.2 Schedule and Morningstar.

Table 4: Bank Delinquency Model

	100×Delinquent (Call Definition)	100×Delinquent (Inc. maturity default and liquidation)	Year-ahead PD (%)
	(1)	(2)	(3)
ln(Balance at Orig.)	0.16**	0.25**	0.37**
	(0.04)	(0.07)	(0.06)
CBD	0.34	0.37	1.30**
	(0.22)	(0.25)	(0.38)
Teleworkable Share	4.17*	2.83	11.16**
	(1.78)	(2.11)	(2.43)
Office	1.21**	1.28**	1.71**
	(0.23)	(0.24)	(0.28)
× ln(Balance at Orig.)	2.19**	2.36**	2.90**
	(0.40)	(0.42)	(0.45)
× CBD	2.55*	3.35**	2.58*
	(1.12)	(1.23)	(1.20)
×Teleworkable Share	17.11**	15.19**	18.04**
	(5.36)	(5.82)	(6.28)
Retail	0.55**	1.00**	0.74**
	(0.16)	(0.23)	(0.22)
Industrial	0.03	0.24	-0.21
	(0.14)	(0.22)	(0.19)
Hotel	1.68**	2.46**	2.95**
	(0.56)	(0.71)	(0.80)
Intercept	0.36**	0.51**	1.27**
	(0.08)	(0.12)	(0.19)
R_a^2	0.027	0.026	0.063
Observations	46,925	46,925	39,419

Notes: This table presents estimates from the equation:

$$100 \times \text{Delinquent}_{i,23} = \alpha_{p(i)} + \beta'(\text{Office}_i \times X_i) + \gamma' X_i + \varepsilon_i,$$

where $\text{Delinquent}_{i,23}$ is a delinquency measure as of 2023:Q4, $\alpha_{p(i)}$ is a fixed effect for loan i's property type, and X_i is a set of risk factors that are observable both in Y-14 and RCA/CoreLogic data: the logarithm of the at-origination loan balance, an indicator for whether the property is in a CBD, and the share of jobs in i's MSA that are identified as being able to be done at home by Dingel and Neiman (2020). Column (1) predicts delinquency for the sample of loans that are on the balance sheet as of the end of 2023. Column (2) presents estimates using the measure of delinquency from Section 2, which includes liquidated and performing ballooned loans as delinquent and paid-off loans as performing. Column (3) presents equivalent analysis predicting the reported year-ahead probability of default. Standard errors, in parentheses, are clustered by bank-origination year. $^{+}$,*,** indicate significance at 10%, 5%, and 1%, respectively.
Sources: Authors' calculations using the Y-14Q H.2 Schedule.

Table 5: Tree Decomposition

	(1) Pr(Delinquent)	(2) Large Bank Share (%)	(3) Small Bank Share (%)
Non-office	0.54%	79.73	82.53
Small Office	1.39%	5.62	13.58
Large Office, Low Telework	8.41%	11.54	3.42
Large Office, High Telework	22.89%	3.11	0.46
Weighted-average delinquency		2.19%	1.03%

Notes: Column (1) presents the delinquency rates for each leaf in the tree model, while columns (2) and (3) provide the portfolio shares for large and small banks, respectively, in each leaf. The fitted delinquency rate for large and small banks—reported in the last row—is the average of (1), weighted by the portfolio shares in (2) or (3), respectively. Small/Large office loans are defined by an at-origination loan balance below/above $23.3 million and Low/High Telework cities have a telework eligible share of employment below/above 44.4%.
Sources: Authors' calculations using the Y-14Q H.2 Schedule, MSCI Real Capital Analytics, and CoreLogic.

APPENDIX

This document contains the supplementary materials as referenced in the manuscript. In Section A.1, we provide additional details on the data. In Section A.2, we provide additional details on the machine-learning estimates. Section A.3 contains the supplementary figures and tables.

A.1. Additional Data Details

We first provide further detail on how we process and harmonize the loan-level panel data for large banks and CMBS. We then describe the CRE origination data from MSCI RCA and CoreLogic.

A.1.1. Additional Details on the Large Bank and CMBS Panel Data

Reporting of loan balances and collateral values for bank loans can be skewed by loan participations and cross-collateralization. For participations, the reported balance is prorated to the size of the bank's participation. When predicting delinquency, we are interested in the size of the loan itself rather than the individual bank's exposure to the loan. Thus, we scale the loan size up by dividing the reported balance by the share of the loan held by the bank to back out the borrower's actual balance. To avoid double counting loans reported by multiple Y-14 banks, we only include observations in which the bank is the one selling the participation interest. Cross-collateralized loans have collateral double-counted as applying to multiple loans (i.e., property values reflect the aggregate value of a collateral pool, but loan balances only reflect the balance on an individual loan). We therefore adjust property values and LTVs by prorating the portion of the collateral attributable to a given loan.

For CMBS data, we exclude agency CMBS loans and defeased loans. Similar to bank loan participations, CMBS loans are often split over multiple pari-passu pieces. To back out appropriate loan sizes, we compute outstanding values by summing balances across such loans, and to avoid double counting, we drop duplicated observations from the sample. Data are updated monthly but not necessarily at month-end. The end of 2022 and end of 2023 sample is as of the last observation reported for those years (i.e., the data from the December data update). When analyzing maturity outcomes, we omit a small number of loans that mature between their December reporting date and year-end.

Some noncore property types have inconsistent identifiers across the bank and CMBS data, so we restrict our attention to core loan categories: multifamily, office, retail, industrial, and hotels. In both data sets, we omit a small number of observations with missing information on geography, property values, or loan balances. We also omit observations with at-origination LTVs that are not between 0 and 0.99 in order to reduce the effect of potential reporting errors. The teleworkable share is missing for loans against properties outside of cities. To avoid systematically dropping these observations, we set the teleworkable share to the value for the 10th percentile of the non-missing sample (about 0.32), under the assumption that these more rural locations are toward the bottom of the telework-exposure distribution. When we add a dummy variable to the OLS specification for whether the teleworkable share is missing, the estimates are near 0, which validates that loan performance is indeed similar for observations without data and observations around that portion of the distribution. In the tree-estimates, this decision means that the low telework split pools markets with a missing or a low telework-eligible share.

A.1.2. CRE Origination Data from Real Capital Analytics and CoreLogic

While detailed, panel data on CRE loans are not generally available for small banks, some information on loan terms at origination are available. This subsection describes how we use RCA and CoreLogic data to form data approximating the composition of at-origination characteristics of outstanding loans across lender types.

We use data from RCA to provide information on CRE loan portfolios for loans with an at-origination balance of $2.5 million or more. While the data do not specify whether mortgages are still outstanding, we can reasonably infer whether they are based on the presence of subsequent transactions. Specifically, we drop loans against properties that are later refinanced or sold, unless the sale is marked as involving the assumption of existing debt. We also drop loans with a maturity date before 2023 in case those loans were paid off without mortgage financing or the refinance does not appear in the data. Loans without a maturity date listed are assumed to have a 10-year loan term. We restrict the sample to loans that finance an already-built investment property to remove types of loans that are not typically part of CMBS deals (i.e., we exclude owner-occupied properties or properties purchased for construction or redevelopment). Data are reported at the property level with loan balances allocated across the properties covered in a deal. To aggregate to the deal level, we sum loan balances within a particular lender-deal id combination. Property types and locations pertain to the largest property in the transaction (by price).

We use data from CoreLogic to provide information on CRE loans with an at-origination value under \$2.5 million. As the sample covers commercial mortgages with open liens, there is no need to impute whether loans are still outstanding. We omit loans flagged as construction loans or owner-occupied loans. Data are reported at the parcel level, with mortgage information repeated when multiple parcels are covered by the lien. To avoid double counting, we only keep the mortgage information associated with the largest property (the parcel with the highest assessed value).[10] Necessary data are sometimes missing, most commonly because a generic "commercial" property type is identified rather than the specific type (e.g., "office," "retail," etc). When computing the portfolio aggregations, we assume that missing observations are reflective of other loans under \$2.5 million and scale up the aggregation weights for the observations where data are available to account for these loans.

Both CoreLogic and RCA report lender names rather than typical bank identifier codes (e.g., ID RSSD). To identify the type of entity making the loan, we fuzzy name match the lender names in each data set to National Information Center institution data. After cleaning to standardize punctuation and other common words (e.g., replacing "national association" with "na" and terms

[10] We assume two parcels are reporting the same mortgage when they have the same mortgage transaction id, lender, loan amount, mortgage date, and mortgage purpose recorded.

such as “bk,” “bancshares,” etc. with ”bank”), we match lender names based on cosine similarity. As some small banks have duplicate names but tend to have a limited geographic footprint, we disambiguate banks with similar names by giving priority to name matches to locally-operating banks (i.e., banks with a branch in the county in Summary of Deposits data). Once a mortgage is matched to a bank, we assess whether it is a large or small bank based on whether the regulatory high holder (if applicable) has more than $100 billion in assets as of 2023:Q4 Call Report/Y-9C data.[11] CMBS loans, which are frequently originated by banks, are identified by whether RCA codes the lender group as CMBS. Any mortgage in CoreLogic that links to a bank is assumed to be a bank portfolio loan, since CMBS generally do not originate loans in the size range for which we use CoreLogic data (see Glancy et al., 2022a).

[11]Loans from foreign banks that have an intermediate holding company subject to stress tests are identified as belonging to a large bank (e.g., loans marked as provided by RBC are attributed to RBC US and counted under the large bank category). Loans from other foreign banks are excluded from the analysis.

A.2. *Machine-Learning Estimators and Hyperparameter Tuning*

The decision tree, KNN, and random forest models are estimated using the DecisionTreeRegressor, KNeighborsRegressor, and RandomForestRegressor classes, respectively, of the scikit-learn package (in Python). Hyperparameters are set to the values that provide the best performance (the estimated default probabilities with the lowest mean squared error) using fivefold cross-validation, stratifying by the outcome variable. Any parameters not discussed here are set to default values.

For the decision tree and random forest estimates, the parameter considered is the "min_impurity_decrease," which determines how much the model needs to improve in order to generate an additional split to the feature space. We search for the best estimator over a grid $[1,2,\ldots 10] \times 10^{-5}$ and find it to be 3×10^{-5} for each estimator. For the tree estimate, we obtain a simple structure with only four terminal nodes and each additional split occurring in the node with the highest probability of default. The tree thus identifies a hierarchy of compounding risk factors: office loans underperform other loans, large loan sizes compound the risk of office loans, and high-telework exposure compounds risks to large office loans.

For the KNN estimates, the parameters considered are the number of neighbors ($K = [100, 200, \ldots 500]$) and the intensity with which we downweight neighbors that are further away (weights decay exponentially in the Euclidean distance between the X-vectors at a rate in $[0,5,\ldots 25]$). We find the optimal parameter values on that grid are 500 neighbors, with weights declining at a rate of 15. Though the number of neighbors is at the top of the grid (suggesting that there could be a benefit to including more neighbors), the gradient from increasing K is flat so we keep K at 500. As this distance and the selection of neighbors is sensitive to the scale of the features, we use min-max scaling for any continuous feature.

A.3. Supplementary Tables and Figures

Figure A1: Decision Tree Delinquency Estimates

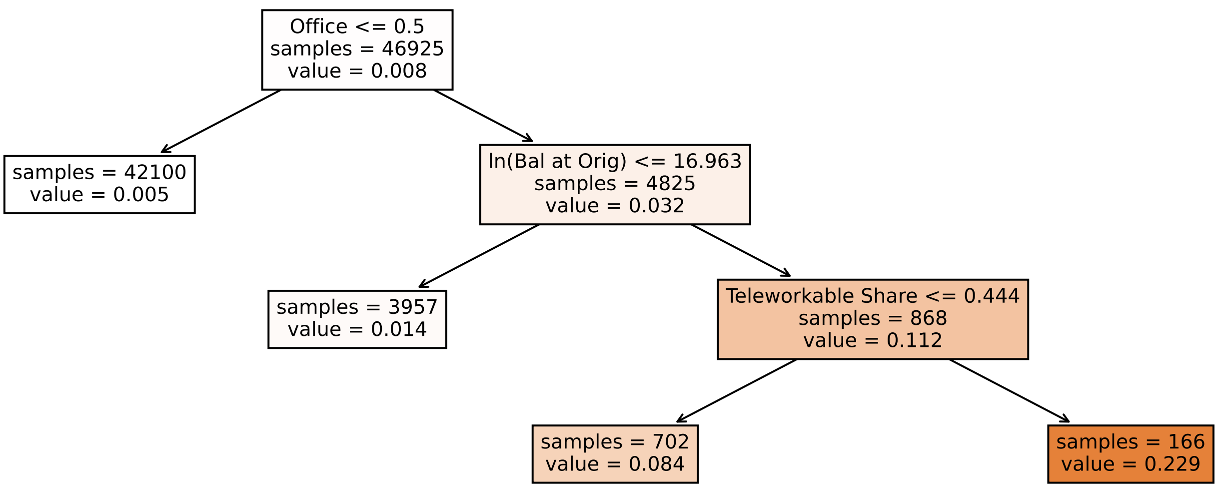

Notes: Each node defines the split that occurs at the node (if there is one), the number of observations in the training data in that node (samples), and the share of those observations that are delinquent (value). Nodes to the left correspond to feature values under the splitting threshold, and nodes to the right are above the threshold.
Sources: Authors' calculations using the Y-14Q H.2 Schedule.

Figure A2: Correlation across Model Predictions

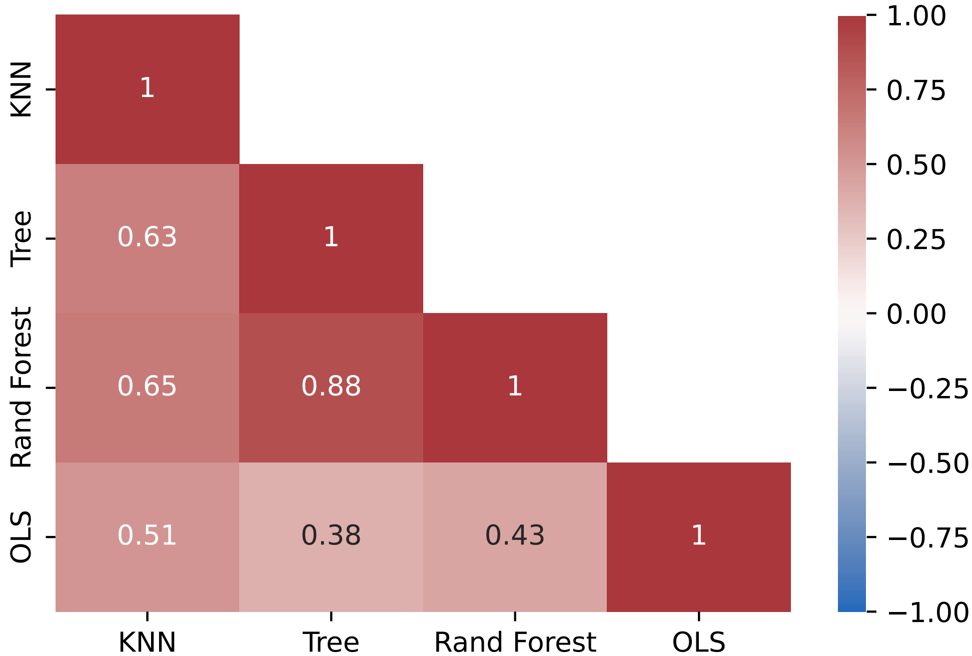

Notes: Correlations of fitted delinquency rates in the RCA/CoreLogic sample.
Sources: Authors' calculations using the Y-14Q H.2 Schedule, MSCI RCA, and CoreLogic.