Determinants of Recent CRE Distress: Implications for the Banking Sector *

David Glancy^{†1}and Robert Kurtzman^{‡1}

¹Federal Reserve Board

August 26, 2024

Abstract

Rising interest rates and structural shifts in the demand for space have strained CRE markets and prompted concern about contagion to the largest CRE debt holder: banks. We use confidential loan-level data on bank CRE portfolios to examine banks' exposure to at-risk CRE loans. We investigate (1) what loan characteristics are associated with delinquency and (2) to what extent the portfolio composition of major CRE lenders determines their exposure to losses. Higher LTVs, larger property sizes, and greater local remote work tendencies are all associated with increased delinquency risk, particularly for office loans. We use several machine learning algorithms to demonstrate that variation in exposure to these risk factors can account for most of the performance disparity across different types of CRE lenders. The headline result is that small banks' comparatively modest delinquency rates mostly reflect observable portfolio characteristics—predominantly their low holdings of large-sized office loans—rather than unobserved factors like extension or modification tendencies.

Keywords: commercial real estate, banks, CMBS

JEL Classification: G21, G22, G23, R33

^{*}We thank Joe Nichols and seminar participants at the Federal Reserve Board R&S workshop 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 (Board of Governors of the Federal Reserve System, 2023). As CRE is the largest loan category on banks' books, these developments have caused concern about CRE loan losses exacerbating other recent banking sector strains (Acharya et al., 2023).¹ Analyzing such effects is complicated by a lack of detailed data on banks' CRE loan holdings. Due to this data limitation, researchers have mostly relied on aggregate bank portfolio data (Faria-e Castro and Jordan-Wood, 2023, 2024), or data from CRE segments with more public reporting (Gupta et al., 2022; Jiang et al., 2023; Glancy and Wang, 2023) to assess risks posed to the banking sector. However, banks serve a selected segment of the CRE market (Glancy et al., 2022a), and recent CRE stresses have been highly uneven (Marsh and Pandolfo, 2024), so extrapolating across different parts of the CRE market can be difficult.

Consistent with this segmentation, Figure 1 shows that loan performance differs markedly across different types of CRE lenders. Though nonperforming loan (NPL) rates for bank and commercial mortgage-back securities (CMBS) loans were comparable during the Global Financial Crisis (GFC) of 2007–09, periods of strain since then have been more confined to CMBS markets; CMBS delinquencies (red) rose moderately in 2016 and spiked at the onset of the COVID-19 pandemic, while delinquencies at large (blue) and small (green) banks remained under 2 percent during these episodes.²

In 2023, delinquency rates for CMBS and large banks rose by similar amounts, but the performance of CRE loans at smaller banks remained strong. CRE strains appear to be highly concentrated: the rise in bank CRE NPLs is driven by nonowner-occupied, nonresidential properties at large banks

¹As of March 6, 2024, total outstanding CRE loans were estimated to be about \$3.0 trillion, compared to \$2.8, \$2.6 and \$1.9 trillion for commercial and industrial, residential real estate, and consumer loans, respectively, according to data from the H.8 release of the Board of the Governors of the Federal Reserve System.

²Even during the GFC, NPL rates for loans secured by existing properties rose less at banks than for CMBS. However, bank NPL rates were pulled up by a near 20 percent delinquency rate for construction and land development loans (a segment not served by CMBS). See Appendix Figure B.1 for a decomposition of bank NPL rates by CRE subcategory.



Figure 1: Nonperforming Loan Rates over Time

Notes: The figure 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. Calculations exclude defeased and real estate owned loans.

Sources: Call Reports, Morningstar, and authors' calculations.

(see Figure B.2), and within this segment, the increase is almost wholly attributable to office loans (see Figure B.3). This heterogeneity highlights that CRE cannot be viewed as a single strained asset class. Understanding the implications of the recent turmoil requires a detailed understanding of where loan performance is deteriorating and which lenders are exposed to those troubled segments.

In this paper, we compile data from a variety of sources to analyze why CRE loan performance differs across lenders. The analysis proceeds in two steps. First, we use loan-level panel data on CRE loan performance to investigate what factors are associated with higher delinquency rates and why loan performance differs between banks and CMBS. To accomplish this, 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 nearly 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

Figure 2: CRE Exposure by Bank Size



Notes: The figure plots CRE as a share of assets (left) and CRE delinquency (right) by bank size. CRE includes nonfarm nonresidential, 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:* Call Reports and authors' calculations.

is smaller and not as clearly attributable to other observable characteristics.³

In the second part of the paper, we investigate potential drivers of the strong performance at small banks. Studying the performance of CRE loans at large banks (i.e., banks with over \$100 billion in assets) is useful as (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. Indeed, Figure 2 shows that banks with between \$1 billion and \$10 billion in assets tend to have the highest concentrations in CRE (left), but also had quite low CRE delinquency rates as of end of 2023 (right). These results indicate that the deterioration in loan performance experienced *to date* is unlikely to cause significant problems for the banking sector. However, 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.

While detailed loan-level panel data are generally not available for small banks, some broad char-

³We provide some suggestive evidence that a greater tendency to modify loans may contribute to these unobserved differences.

acteristics pertaining to loan sizes, property types, and locations are available due to the public reporting of mortgage liens. We can therefore estimate models of loan performance using data from large banks and investigate the degree to which these observable factors can explain the strong performance at small banks. We do so using a variety of models in order to provide predictions from both easily interpreted models where it is clear what variables drive performance differences (OLS and low-complexity trees) and more flexible ones that may better reflect complex patterns in the data (K-nearest neighbors and random forests).

Regardless of the model chosen, we find that differences in performance of CRE loans at large and small banks can mostly be attributed to differences in loans sizes, and in the types and locations of the properties securing the loans. For the models that allow us to meaningfully decompose differences in loan performance between large and small banks, 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 NPL rates.

Overall, these findings demonstrate that the resiliency of small banks' CRE portfolios can be rationalized with property-level observables. If NPLs at small banks were low because their high concentrations in CRE loans prompted them to delay loss recognition by "evergreening," such differences would not be picked up by our model. Instead, it appears that small banks' NPLs are mostly lower because their loans are safer on observed dimensions. Therefore, the primary risk to small banks comes less from the CRE strains experienced through 2023, but more from the risk of stress spreading to other parts of the CRE market—for example due to an economic downturn or higher-for-longer interest rates further pressuring valuations.

The rest of the paper proceeds as follows. Section 2 reviews an array of factors identified in the literature that could contribute to differences in loan performance across CRE lenders. Section 3 identifies what factors are associated with delinquency at large banks and CMBS. Section 4 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 5 concludes.

2. LITERATURE ON DIFFERENCES ACROSS CRE LENDERS

To guide our analysis of the factors potentially contributing to the differences in loan performance across lenders, we first review the literature pertinent to this question. We split this discussion into two parts, separately discussing the factors that could mitigate or amplify loss for bank CRE loans.

2.1. Factors That Can Support Bank Loan Performance

Modification Ability The first factor supporting bank loan performance is that banks are more able to modify distressed loans to preempt foreclosure (Black et al., 2017, 2020). Indeed, Glancy et al. (2022b) present evidence that forbearance policies supported bank CRE performance at the onset of the pandemic. Bank regulators released a policy statement on CRE loan workouts in 2023 reaffirming that "prudent CRE loan accommodations and workouts are often in the best interest of the financial institution and the borrower", thus demonstrating that support for bank workouts is more than a COVID-era phenomenon.⁴ In contrast, CMBS face limitations to modifying loans due to IRS policies, stipulations in pooling and servicing agreements, and conflicts of interest across disparate investors (Wong, 2018; Flynn Jr. et al., 2023). While the nature of workouts will presumably be different in the current environment than at the onset of the pandemic (e.g., extensions rather than forbearance), banks' flexibility in this regard may enable them to prevent or more quickly resolve loan stress.

⁴The guidance is available on the Federal Reserve's website at https://www.federalreserve.gov/supervisionreg/srletters/SR2305a1.pdf.

Recourse and Lower Leverage A second factor supporting the performance of bank CRE loans is that borrowers typically have more equity at stake. First, bank loans typically have lower LTVs, which gives more room for property values to fall before the lender starts taking losses (Glancy et al., 2022a). Second, bank CRE loans are predominantly recourse loans, meaning that borrowers have assets besides the subject property at stake if they default. Glancy et al. (2023) show that most CRE loans from large U.S. banks have recourse. The authors demonstrate that these recourse loans receive more favorable underwriting terms and were less likely to require forbearance at the onset of the pandemic, both consistent with recourse supporting loan performance. Unlike bank loans, CMBS loans are generally non-recourse (barring springing guarantees for particular "bad acts"). As such, CMBS borrowers may be more willing to walk away from a property even if they have the ability to service the debt on it.⁵

Differences in Loan Size/Location A third factor potentially contributing to differences in loan performance is that banks tend to lend against smaller properties than CMBS (Ghent and Valkanov, 2016; Glancy et al., 2022a).⁶ While loans against larger properties are not considered inherently riskier, they have underperformed in the current environment. Part of this effect likely reflects location; high priced properties are disproportionately located in central business districts (CBDs). Amid the shift to working from home, CBDs have experienced declines in commuting activity (Monte et al., 2023) and commercial rents (Rosenthal et al., 2022), and a greater deterioration in property occupancy and income following lease expirations (Glancy and Wang, 2023). Relative to CMBS, office loans at banks, and especially small banks, are more likely to be in suburban markets, which have been less affected by the shift to remote work.⁷ Even within a given location, larger offices likely have a tenant mix placing them at a greater risk of departures (e.g., more space

⁵Recourse might be particularly valuable to lenders in the current environment given that stress so far has been largely confined to the office sector. If sponsors or guarantors have other assets of value—for example, houses, equities, or less-troubled CRE properties—owners may be more willing and able to maintain payments on loans secured by troubled properties.

⁶By virtue of their diversified customer base, CMBS are able to fund larger loans that would generate too much concentration risk for a balance sheet lender.

⁷Glancy and Wang (2023) show that small and regional banks have a lower exposure to office loans in CBDs or markets with a greater increase in remote work.

accounted for by multi-location tech companies).

Better Screening The last factor we highlight supporting the performance of bank loans is that banks typically retain the credit risk for the loans they origination, and thus can have better screening incentives. Work on the CMBS market before the GFC provides evidence that adverse selection (An et al., 2011) and moral hazard (Ashcraft et al., 2019) reduce the quality of loans in CMBS pools. In response, as part of the Dodd-Frank Act, regulators implemented risk retention requirements to reduce such agency problems (Flynn Jr et al., 2020). Nonetheless, Griffin and Priest (2023) present evidence that agency conflicts between originators and CMBS investors persist and may create underwriting weaknesses that are revealed during times of stress.

2.2. Factors That Can Hinder Bank Loan Performance

Specialization in (Historically) Riskier Loans While greater renegotiation flexibility may support the performance of any given bank loan, this effect may give banks an advantage in financing riskier properties where such flexibility is more valuable (Black et al., 2017, 2020). One clear dimension along which risk differs is that CMBS lend against income-producing properties, whereas banks originate more bridge and construction loans. The performance of bridge and construction loans can be more dependent on the particular business model of the borrower and thus subject banks to the risk that properties are more difficult to lease-up than expected or fail to earn returns sufficient to recoup the alteration costs. Consistent with these loans being riskier, construction and land development loans are subject to higher FDIC fee assessments, receive higher capital requirements when LTV limits are not met (Glancy and Kurtzman, 2022), and disproportionately contributed to bank failures during the GFC (Friend et al., 2013).⁸

⁸For more details on these loans receiving higher FDIC fee assessments, see FDIC Law, Regulations, Related Acts, Title 12 Chapter III Subchapter B Appendix C to Subpart A to Part 327 available at https://www.ecfr.gov/current/title-12/chapter-III/subchapter-B/part-327.

More Financially Constrained Borrowers A final difference between bank and CMBS loans is that banks may cater to more financially constrained borrowers. Roughly half of CMBS lending goes to public or institutional buyers, compared to only 10 percent of bank lending (Glancy et al., 2022b). These larger sponsors with more diversified funding sources can potentially maintain loan payments in the face of a disruption to property cash flows (assuming they find it optimal to do so). Similar selection effects occur outside the CRE market, where banks disproportionately serve smaller, younger, and riskier firms, while more established firms utilize market financing (Petersen and Rajan, 1994; Bolton and Freixas, 2000).

3. CRE LOAN PERFORMANCE AT LARGE BANKS AND CMBS

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

3.1. Methodology

As the last section discusses, there are a host of factors that may contribute to differences in the performance CRE loans across lenders. We now empirically examine how much these factors matter in the current environment. 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 over \$100 billion in assets. The CMBS data come from Morningstar, which compiles loan-month data from CMBS disclosures.⁹ We classify lenders by who holds the loan, rather than who originates it, so bank-originated loans in CMBS pools are considered CMBS loans. More information on how we clean and harmonize these data are in Section A.1.

Though banks below the Y-14 reporting threshold tend to have higher concentrations in CRE, an advantage of the Y-14 sample is that it covers the group of banks for which CRE loan performance

⁹We 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.

has materially deteriorated. This data is thus useful for examining the factors causing CRE loans to go delinquent. While we cannot analyze the performance of CRE loans at smaller banks in such detail, we can assess the performance of Y-14 loans with characteristics resembling those of small banks (i.e., smaller loans in smaller markets) to explore the implications for smaller banks. We conduct this exercise in Section 4.

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

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

for the sample of bank or CMBS loans that were outstanding as of the end of 2022. Delinquent_{*i*,23} is in indicator for whether loan *i* is delinquent as of the last observation in 2023, defined here as being past due, performing beyond its maturity date, or liquidated.¹⁰ 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. The coefficient on CMBS reflects the difference in delinquency rates for CMBS loans compared to those from large domestic banks, controlling for the two characteristics most strongly related to loan performance: whether they matured (capturing difficulty refinancing in an environment of falling valuations and tight lending standards) and whether the loan is secured by an office property (the property type accounting for most of the rise in delinquencies over 2023).

 $X_{i,23}$ is a set of controls that includes other property type dummies.¹¹ In some specifications, we layer in additional controls to assess whether they can account for some of the difference in delinquency between large banks and CMBS. These additional controls include LTV, property size, a recourse indicator, and geographic characteristics (whether the property is in a CBD or city where

¹⁰The quarter of observations is either 2023:Q4 if the loan is active as of the end of the year or the quarter the loan paid-off or was liquidated otherwise.

¹¹These variables are included in every specification but coefficients not displayed. Multifamily is the omitted category. Hotel and retail are the next-most-likely loans to go delinquent after office, while industrial delinquency rates are not significantly different from those for multifamily loans.

more jobs can be done remotely). Observing how β_1 changes as extra controls are added provides information on the extent to which these characteristics can account for differences in loan performance. In the most expansive specifications, we also include controls for the occupancy of the property and debt yield of the loan.¹² These controls thus assess the extent to which differences are driven by unobserved risks to property performance that materialize as a deterioration in financial performance. Finally, in robustness exercises we conduct similar analysis predicting whether loans receive extensions to assess whether banks are providing more accommodation to stressed borrowers.

Summary statistics of the main variables of interest are shown in Appendix Table B.1, with data on large bank 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 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 one percentage point when using a narrower definition that only counts past-due or nonaccrual loans as delinquent (which is more aligned with the NPL measures from the Call Report). Columns (3) and (4) present the same data, but weighted by loan size, thus making the averages more reflective of aggregate portfolio shares. Broad patterns pertaining to differences between bank and CMBS portfolios are similar, but delinquency rates are higher, reflecting the fact that performance has deteriorated more for larger loans.

¹²Debt yield is the ratio of net operating income to the outstanding loan balance and thus reflects the ability of a property's cash flows to pay off the loan. Since 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 prior.

3.2. Results

The main estimates are reported in Table 1. Column (1) presents estimates from the most parsimonious specification, which only includes the CMBS, maturing loan, and property type indicators. The results show that 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 loans that mature 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 the loan off as it comes due.

To investigate why CMBS loans have higher delinquency rates, we incorporate different factors discussed in the previous section as potentially contributing to differences in loan performance across lenders. Column (2) adds in the at-origination LTV of the loan, the logarithm of the at-origination property value, and an indicator for whether the loan has recourse. The size and LTV controls account for the fact that CMBS provide higher leverage, on average, and tend to lend against larger properties. Since CMBS loans are essentially entirely non-recourse, the effect of recourse is identified off of differences in the performance of bank-held recourse and non-recourse loans. Therefore, the coefficient on CMBS is now estimated off of the difference in the performance between CMBS loans and non-recourse bank loans.

The findings show that the characteristics associated with bank loans—lower leverage, smaller properties, and recourse—are also associated with stronger loan performance. A one standard deviation higher LTV (0.16) or property size (1.16) are associated with delinquency rates that are roughly 80 and 100 basis points higher, respectively.¹³ The effect of recourse is smaller, with recourse loans having a delinquency rate that is about 40 basis points lower than similar non-recourse loans. Adding these three additional controls reduces the coefficient on CMBS from 1.65

¹³Variation in property values reflect differences in both size (i.e., square footage) and valuation (i.e., price per square foot). Most of the effects on delinquency are attributable to differences in size, so we refer to the variable ln(Value at Orig.) as "size" for the sake of brevity.

to 0.45, indicating that CMBS' tendency to make larger, higher LTV, non-recourse loans accounts for about two-thirds of their inferior performance.

Column (3) adds in geographic characteristics of the property securing the loan: whether the property is in a CBD and the share of jobs in the metropolitan statistical area (MSA) that are identified as being able to be done at home by Dingel and Neiman (2020).¹⁴ The results indicate that delinquency rates are about 2 percentage points higher in CBDs, but only modestly higher for cities more exposed to a potential shift to remote work; a one standard deviation increase in the teleworkable share (0.046) increases delinquency by about 25 basis points. Adding the extra geographic controls causes the estimated coefficient on CMBS to increase a bit; though banks generally do less lending in adversely affected markets, this is mostly due to size differences which were accounted for by the specification in column (2). Adding these variables also only reduces the coefficient estimate on ln(Value at Orig.) very slightly, indicating that the effects of property size are not driven by these geographic risk factors. In fact, most estimates are little changed with the inclusion of ZIP code fixed effects (not shown), so the risks associated with larger loans appears to predominantly reflect non-geographic factors.

Column (4) adds in two variables pertaining to the property's financial situation: the occupancy and an indicator for whether the property's debt yield (net income as a share of the loan balance) is under 8 percent. These controls thus account for whether differences in performance are driven by variation in occupancy or cash flow risks across lenders. Low occupancy or debt yields are highly predictive of delinquency: a 10 percentage point drop in occupancy raises the probability of delinquency by 1.5 percentage points, and a low debt yield raises the probability of delinquency by about 1.9 percentage points.

While the overperformance of bank loans can be partially explained by observable characteristics, there remains a large unobserved component. CMBS loans are about 0.9 percentage points more likely to go delinquent than bank loans with similar underwriting terms and even similar financial

¹⁴We identify a property as being in a CBD if its ZIP code belongs to a submarket that CBRE defines as being in a CBD.

	$100 \times \text{Delinquent}_{i 23}$							
	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_a^2	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?	<u>√</u>	<u>√</u>	<u> </u>	<u> </u>				

Table 1: Loan Performance by Lender Type

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 under 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: Y-14Q H.2 Schedule, Morningstar, and authors' calculations.

performance. One explanation for this result is that banks are more able to renegotiate stressed loans. Some suggestive evidence of this can be found in Appendix Table B.2, which shows that bank loans were about 5 percentage points more likely to receive extensions compared to CMBS loans. Without a counterfactual saying what would have happened to these loans if they did not receive an extension, it is difficult to say precisely how banks' willingness to extend loans affected performance differences. However, if some of these extended loans would have otherwise gone delinquent, such accommodation could have contributed to the comparative strength of bank loans.

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

Column (6) shows that the underperformance of CMBS office loans can be entirely attributed to differences in loan sizes, LTVs, and the use of recourse. Office loans have delinquency rates that are 3.6 percentage points lower when they have recourse, 1.3 percentage points higher when the LTV is 10 percentage points higher, and about 2.3 percentage points higher for properties with an at-origination property value that is one standard deviation higher. Once these factors are controlled for, CMBS loans are found to have delinquency rates that are slightly *lower* than bank loans. Put differently, CMBS office loans perform somewhat worse than nonrecourse bank loans of similar size and LTV. However, CMBS have higher delinquency rates overall due to differences in these three characteristics.

Column (7) adds in geographic controls, which have a much larger effect on office loans than for CRE properties in general: the effect of being in a CBD and telework exposure are about doubled for office loans relative to the overall CRE loan pool. These controls reduce the coefficient on ln(Value at Orig.) somewhat, but most of the effect of size is not driven by these geographic characteristics.

Finally, column (8) adds in the controls for occupancy and debt yield. Weaker financial perfor-



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: Y-14Q H.2 Schedule, Morningstar, and authors' calculations.

mance are more predictive of delinquency within the office sector. Including these controls reduces the predicted effects of size and the geographic risk factors, indicating that the higher delinquency rates for loans with those characteristics are partly attributable to worse declines in occupancy or income. The fact that those variables remain significant even in the presence of the additional financial controls likely reflects the fact that property performance is slow moving. Owners of larger properties in more adversely affected markets may go delinquent in anticipation of future property strains, especially if those anticipated strains leave them unable to refinance upon maturity.

To recap the findings, many of the drivers of CRE delinquencies are as would be predicted given the nature of the stress. The largest driver is loan maturity, reflecting problems refinancing balloon payments amid tighter credit conditions, lower valuations, and higher interest rates. Additionally, delinquency is higher for properties most exposed to the decline in demand for space—namely, office properties and properties in CBDs or markets more exposed to remote work.

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—such as a greater ability to modify loan terms, or borrowers' concerns about damaging existing bank relationships—are supporting overall loan performance.

Second, loans securing large properties have higher delinquency rates, suggesting that they have unobserved risk characteristics—for example, a worse income outlook, more difficulty making up operating shortfalls, or sponsors more willing to strategically default—that are weighing on loan performance. The magnitude of this effect can be most clearly seen in Figure 3, which presents binscatter estimates of how loan delinquency varies by property size and lender type. The delinquency rate for bank CRE loans is modest for properties with an at-origination value under \$20 million, even for office loans. Delinquency is instead concentrated in large CMBS-funded properties, or large bank-funded offices. The next section explores the extent to which such differences across property sizes can explain the stronger loan performance at smaller banks.

4. IMPLICATIONS FOR SMALL BANKS

This section explores the implications of these findings for the performance of CRE loans at small banks. To the extent that small banks lend against even smaller and less urban properties, the patterns documented in Section 3 may contribute to their comparatively strong loan performance. To test this, we first compile data from various sources on the composition of CRE loan portfolios across different types of lenders. We then estimate a set of models of bank loan performance using only the characteristics that are observable for small banks' CRE loan portfolios. Lastly, we examine the fitted delinquency rates to assess the extent to which those factors can account for small banks' lower delinquency rates.

4.1. 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). RCA 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 dollars 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 for loans below this threshold.

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 CRE loans). This sample covers the CRE loans that have exhibited the most stress through 2023 (see Appendix Figure B.2). More information on how we construct the data, including details lender name matching, sample selection, and how we impute whether RCA loans are still outstanding are in Section A.2.

After harmonizing, name matching, and combining these data sources, we have a cross-section of at-origination characteristics for what should be the near universe of outstanding commercial mortgages and identifiers for the type of lender that originated the loan. For each of these loans *i* from one of the lender types *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 3, 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(Balance at Orig.) in place of the measures of LTV and property value. Since 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 based on this data are shown in the last two columns of Table B.1. The averages for large banks are broadly in line with the weighted averages of the Y-14 data shown in column (3), suggesting that the data we construct from RCA/CoreLogic successfully matches banks' true portfolio composition. The statistics shows 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.

4.2. Bank Delinquency Models

The second step to 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 estimating $\hat{D}(\cdot)$. In order to be able to benchmark the fitted delinquency rates against those observed in the data, we use a sample of loans and definition of delinquency aligned with the non-performing loan rates available in the Call Reports. 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+ days past due or nonaccrual.

In choosing a model to estimate the probability of delinquency, we face a trade-off 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 as to 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 sample by feature values so as to achieve the best model improvement with each split. We set a high threshold for splitting to reduce the number of leafs and simplify interpretation. The algorithm thus splits the feature space into a small number of regions, and the predicted delinquency rate is the share of Y-14 loans within the region that are delinquent.
- 3. **K-nearest neighbors (KNN)**: Finds the observations in the Y-14 data with the most similar feature values, 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**: Repeatedly takes bootstrapped samples of the training set and fits shallow trees to the samples. The estimated delinquency rate is an average of the prediction of the trees across the samples.

For the OLS estimator, the primary decision is the specification. Motivated by the results in Table 1, which show that the main risk factors considered disproportionately affect office loan performance, we choose a mostly linear model, with an 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 hyperparameters. For each model, we search over a parameter grid and use stratified 5-fold cross validation to choose the parameters that produce the lowest mean-squared error in the left-out data. More detail on these estimators and hyperparameter tuning are in Section A.3.

The coefficient estimates for the linear model are in column (1) of Table 2. Variables pertaining to loan size and telework ability are demeaned so the coefficient on office shows the higher delinquency rate for an office property outside of a CBD with an average level for other risk factors.

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 higher LTV loans. We also see higher delinquency rates in CBDs and cities where more jobs can be done remotely, particularly for office properties.

Column (2) presents estimates from the same specification but with the broader sample and definition of delinquency used in the previous analysis. The main findings generally hold, but the property-specific intercepts are a bit higher. Namely, the predicted delinquency rates are somewhat higher since they better account for maturity-defaults, but the drivers of differences in delinquency are not meaningfully different.

Columns (3) supplements the predictions of delinquency with a prediction of the year-ahead probability of default based on banks' internal risk ratings. If small banks are only outperforming larger banks because they have loans that are expected to deteriorate later (e.g., if stress were to start in CBDs before spreading to other markets), these forward-looking measures would allow us to pick up whether smaller banks have characteristics associated with an expected *future* deterioration in performance. For the most part, the factors associated with delinquency tend to be associated with expected future delinquency with a broadly similar intensity. If anything, estimated effects in column (3) tend to be a bit higher than estimates in columns (1) and (2), indicating that banks expect the risk factors associated with delinquency so far to be associated with somewhat further deterioration in performance going forward.

The delinquency rate estimates for the decision tree estimator are visualized in Figure B.4. The first split in the tree is by whether the loan is secured by an office, the second is by whether the

	100×	Year-ahead	
	(Call Definition) (Inc. maturity default and liquidation)		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)
\times ln(Balance at Orig.)	2.19**	2.36**	2.90**
	(0.40)	(0.42)	(0.45)
imes CBD	2.55*	3.35**	2.58*
	(1.12)	(1.23)	(1.20)
\times 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

Table 2: Bank Delinquency Model

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 Delinquent_{*i*,23} is a delinquency measure as of 2023:Q4, $\alpha_{p(i)}$ is a fixed effect for *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 3 (which includes liquidated and performing ballooned loans as delinquent and paid-off loans as performing), and columns (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: Y-14Q H.2 Schedule and authors' calculations.

office loan is large-sized (an at-origination balance over \$23.3 million), and the third is based on whether the large office loans are in a high-telework-eligible market (Teleworkable Share> 0.44). The delinquency rates in the decision tree estimator are 0.5% for non-office loans, 1.4% for small office loans, 8.4% for large office loans in low-telework markets, and 23% for large office loans in high-telework markets.

The KNN and random forest estimators do not have as clear a correspondence between features and the probability of 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 Figure B.5). 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.

4.3. Cross-bank Differences in Fitted NPLs

What do these estimates imply for differences in loan performance across lenders? We now use the delinquency models discussed in Section 4.2 to generate predicted delinquency rates for different types of lenders using the cross-lender portfolio data discussed in Section 4.1. The object of interest is the expected delinquency rate: Fitted Delinquency $_{j} = \sum_{i|j} \omega_{i,j} \hat{D}_{m}(X_{i,j})$, where $\omega_{i,j}$ is the share of *j*'s loan portfolio accounted for by loan *i*. This estimate shows the extent to which differences in the performance of loans across lenders can be attributed to broad differences in the composition of their CRE portfolios. If the drivers of CRE performance at small and large banks are similar, but small banks just perform better because loans are safer on the modeled dimensions, then the fitted delinquency rates should match those observed in the data. If small banks' CRE loans perform better to evergreen), then their stronger performance would be for unobserved reasons, and fitted delinquency rates across banks of different sizes would not differ by much.

Figure 4 plots observed NPLs by lender type based on Call Report data (the blue bars), and the fitted NPLs based on the four delinquency models (the bars below the blue bars). For the types of loans included in the previous analysis (multifamily and non-owner-occupied, nonfarm, non-residential loans), small banks have NPL rates of about 0.6 percent, whereas large banks have NPL rates of around 2.6 percent. This 2-percentage-point differential is well explained by then differences in the observable loan and property characteristics included in the delinquency models. Across the four models, the fitted delinquency rate for small banks ranges from 0.99 percent in the random forest model to 1.21 percent in the OLS model. In other words, about 1.4 to 1.6 of the 2-percentage-point difference in NPLs across large and small banks—70 to 80 percent of the gap—can be attributed to differences in the composition of their CRE portfolio along a relatively small number of dimensions (loan size, property type, and geographic exposure to the shift to remote work).

For large banks, the fitted delinquency rates align closely with the observed ones: fitted delinquency rates range from 2.19 percent to 2.63 percent, relative to 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 the estimates are fit to large bank data, it does increase confidence in the methodology. It would be possible for the fitted delinquency rates to deviate from the actual ones due to sampling problems with the RCA/CoreLogic data or residuals that correlated with loan size (since the portfolio aggregations are weighted).¹⁵ That the predictions align with the observed data suggests that these are not major problems.¹⁶

What drives these differences in fitted NPLs? While the KNN and random forest models are com-

¹⁵The first issue would appear 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 **B.1** also indicates that this is not a problem, as portfolios at large banks in the RCA/CoreLogic data match those in the Y-14 data. The second issue could appear if the way the model accounts for size effects is misspecified. For example, the decision tree estimates only reflect whether office loans are above a particular size threshold, whereas Figure 3 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.

¹⁶There is also some evidence in the cross-section that the model works as intended. Figure **B.6** shows that the fitted delinquency rate at the bank level increases roughly one-for-one with banks' realized non-performing loan rates.



Figure 4: Realized vs. Expected Nonperforming Loan Rates

Notes: The top bars shows the 2023:Q4 nonperforming loan rate for banks with under (top set) and over (bottom set) \$100 billion in assets. The remaining bars give weighted average delinquency rates for loans in small and large banks' portfolios based on a linear probability (red), a regression tree (green), a K-nearest neighbors (yellow), and a random forest (purple) model.

Sources: Y-14Q H.2 Schedule, MSCI RCA, CoreLogic, and authors' calculations.

plex enough to make it difficult to attribute differences in predictions to particular features, the OLS- and tree-based predictions can be decomposed to clarify why loans at large and small banks appear to be performing differently. Figure 5 presents a waterfall chart showing the various factors contributing to differences in loan performance between large and small banks. The first red bar shows the residual, or the unexplained amount by which small banks were overperforming large banks. With the OLS model, small banks are expected to have a nonperforming loan rate of 1.2%, but only had a nonperforming loan rate of about 0.6%, meaning about 60 basis points of their overperformance is driven by factors not in the model. The other bars show how much individual variables in the regression in Table 2 contribute to differences in delinquency rates.¹⁷ While small banks appear to benefit from having smaller loans in general and fewer loans in troubled areas (CBDs or areas with more teleworkable jobs), by far the biggest component explaining their overperformance is the size-by-office interaction; the property size-by-office interaction effect accounts for almost 1 percentage point of the 1.6 percentage point difference in fitted NPL rates.

¹⁷Specifically, 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 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 variables in the regression specification.



Figure 5: OLS NPL Decomposition

Notes: Blue bars show the delinquency 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 2 contributes to the difference (except for the first bar which gives the unexplained component).

Sources: Y-14Q H.2 Schedule, MSCI RCA, CoreLogic, and authors' calculations.

Most of the other effects work in the same direction, but nothing else contributes more than 0.28 percentage points to the differences. Those other variables have either too small an effect on delinquency or not large enough differences across bank sizes to contribute as notably to differences in predicted NPL rates. In short, the OLS estimates indicate that about half of small banks' superior CRE performance can be explained by them being less exposed to large office loans.

We get a similar picture when we decompose the differences in fitted delinquency rates using the tree estimates. Table 3 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). The results show that while smaller banks have a modestly smaller office exposure (roughly 17.5%, relative to 20% for large banks), the bigger differentiator is the size of the office loans. Roughly 15% of CRE lending by large banks is against large office loans (offices

	(1)	(2)	(3)
	Pr(Delinquent)	Large Bank	Small Bank
		Share (%)	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%

Table 3: Tree Decomposition

Notes: Each row gives the delinquency rate for a particular leaf in the tree model (column 1), or the share of large or small banks' CRE portfolios composed of loans in that leaf (columns 2 and 3, respectively). The fitted delinquency rate for a particular lender is the average of (1), weighted by the portfolio shares in (2) or (3), reported in the last row. Small/Large office loans are defined by an at-origination loan balance above/below \$23.3 million and Low/High Telework cities have a telework eligible share of employment below/above 44.4%. *Sources:* Y-14Q H.2 Schedule, MSCI RCA, CoreLogic, and authors' calculations.

with an at-origination balance above \$23.3 million), compared to only about 4% of CRE lending at 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 the case of high-telework areas), this composition produces large differences in the fitted nonperforming loan rates of large and small banks' CRE portfolios.

5. 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. For office loans, these differences can be accounted for by banks making smaller loans, where borrowers have more skin in the game (from either recourse or property equity). These factors only partially explain the differences in performance of non-office loans, suggesting that other factors, such as banks' willingness to work out stressed loans, also

contribute. 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).

The findings of this paper suggest that some of the more dire scenarios pertaining bank CRE losses are unlikely to come to fruition. Small banks' comparatively low nonperforming loan rate appears to reflect the composition of their loan portfolio rather than, for example, "extend and pretend" activities delaying the realization of delinquency. Though large banks are more exposed to adversely affected segments of the CRE markets, they generally have high capital levels and low CRE concentrations that should enable them to weather those losses.

All told, these findings suggest that strains would need to expand to other segments of the CRE market to cause systematic problems for the banking sector. Of course, just because small banks' CRE loans have characteristics that have insulated them so far, this does not mean that they will be immune to future stresses. Some features of bank loans-modification ability, lower leverage, and lower at-risk office exposure—are likely to continue supporting performance in the future. However, there are some segments that have performed reasonably well that could reasonably deteriorate in the future. First, the reason that smaller loans have overperformed so much is not clear, which makes it hard to trust that the effect will persist to the same extent. If small office loans start to perform more similarly to large ones, delinquency at small banks would increase notably. Second, while multifamily delinquency rates are low, they are rising amid elevated interest rates, operating expenses and competition from new supply. Third, this analysis has focused predominantly on CRE loans secured by income-producing properties, whereas construction loans were a primary drivers of stress in past crises. While the delinquency rate on bank construction loans has been modest so far, difficulties leasing new space or obtaining stabilized financing could create strains in the future as more projects exit construction into a challenging environment. Thus while this study provides a description 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. P. RICHARDSON, K. L. SCHOENHOLTZ, B. TUCKMAN, R. BERNER,
 S. G. CECCHETTI, S. KIM, S. KIM, T. PHILIPPON, S. G. RYAN, ET AL. (2023): "SVB and beyond: The banking stress of 2023," *Available at SSRN 4513276*.
- AN, X., Y. DENG, AND S. A. GABRIEL (2011): "Asymmetric information, adverse selection, and the pricing of CMBS," *Journal of Financial Economics*, 100, 304–325.
- 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, 5173–5211.
- BOARD OF GOVERNORS OF THE FEDERAL RESERVE SYSTEM (2023): "Financial Stability Report." 16–17. url=https://www.federalreserve.gov/publications/files/financial-stability-report_20230508.pdf.
- BOLTON, P. AND X. FREIXAS (2000): "Equity, Bonds, and Bank Debt: Capital Structure and Financial Market Equilibrium Under Asymmetric Information," *Journal of Political Economy*, 108, 324–351.
- 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," Tech. rep., Federal Reserve Bank of St. Louis.

- FLYNN JR., S. J., A. GHENT, AND A. TCHISTYI (2023): "The Imitation Game: How Encouraging Renegotiation Makes Good Borrowers Bad," *Working paper*.
- FLYNN JR, S. J., A. C. GHENT, AND A. TCHISTYI (2020): "Informational efficiency in securitization after Dodd-Frank," *The Review of Financial Studies*, 33, 5131–5172.
- FRIEND, K., H. GLENOS, AND J. B. NICHOLS (2013): "An Analysis of the Impact of the Commercial Real Estate Concentration Guidance," Tech. rep., Office of the Comptroller of the Currency.
- GHENT, A. C. AND R. VALKANOV (2016): "Comparing securitized and balance sheet loans: Size matters," *Management Science*, 62, 2784–2803.
- 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, 2029–2080.
- GLANCY, D. AND R. KURTZMAN (2022): "How Do Capital Requirements Affect Loan Rates? Evidence from High Volatility Commercial Real Estate," *Review of Corporate Finance Studies*, 11, 88–127.
- 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– 1136.
- 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, 2233–2276.
- GUPTA, A., V. MITTAL, AND S. VAN NIEUWERBURGH (2022): "Work from Home and the Office Real Estate Apocalypse," *NBER Working Paper 30526*.
- JIANG, E. X., G. MATVOS, T. PISKORSKI, AND A. SERU (2023): "Monetary Tightening, Commercial Real Estate Distress, and US Bank Fragility," *NBER Working Paper 31970*.
- MARSH, W. B. AND J. PANDOLFO (2024): "Banks' Commercial Real Estate Risks Are Uneven," *Economic Bulletin*, 1–4.
- MONTE, F., C. PORCHER, AND E. ROSSI-HANSBERG (2023): "Remote Work and City Structure," *American Economic Review*, 113, 939–981.
- PETERSEN, M. A. AND R. G. RAJAN (1994): "The benefits of lending relationships: Evidence from small business data," *The Journal of Finance*, 49, 3–37.
- 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.
- WONG, M. (2018): "CMBS and Conflicts of Interest: Evidence from Ownership Changes for servicers," *The Journal of Finance*, 73, 2425–2458.

A. DATA APPENDIX

A.1. Large Bank and CMBS Panel Data

Loan level panel data are available for CRE loans held by banks with over \$100 billion in assets and for CRE loans held in CMBS pools, but from different data sources. Here we describe in more detail how we process and harmonize these data sources.

Banks serve some CRE-segments that CMBS do not. To focus on areas where both lenders overlap, we restrict the bank sample to first lien, non-owner occupied, non-construction CRE loans. Data is reported as of quarter end, so analysis in Section 3 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 analysis in Section 4 uses the sample of loans were still on balance sheet as of the end of 2023 (to align with Call Report reporting).

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, so 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 where 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 is 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 outomes, we omit a small number of loans that mature between their December reporting date and year end.

Some non-core 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 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 towards 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.

A.2. CRE Origination Data

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 atorigination balance of \$2.5 million or more. While the data does 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 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 is 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 is 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).¹⁸ Necessary data is 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 is available to account for these loans.

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

¹⁸We 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.

"bank"), we match lender names based on the cosine similarity of TFIDF vectors (identifying lenders with similar character bigrams). Since 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 in a county where the bank has a branch (using Summary of Deposits data). Once a mortgage is matched to a bank, we assess whether they are 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.¹⁹ 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 that size range (see Glancy et al., 2022a).

A.3. Machine Learning Estimators and Hyperparameter Tuning

The three machine learning estimators are estimated in Python. The decision tree, K-nearest neighbors and random forest models are estimated using the DecisionTreeRegressor, KNeighborsRegressor, and RandomForestRegressor classes, respectively, of the scikit-learn package. Hyperparameters are set to the values that provide the best performance (the estimated default probabilities with the lowest mean squared error) using 5-fold cross-validation, stratifying by the outcome variable. Any parameters not discussed here are set to default values.

For the decision tree 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 $[1,2,...10] \times 10^{-5}$ and find it to be 3×10^{-5} . This results in a tree with 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

¹⁹Loans 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.

loans.

For the KNN estimates, the parameters considered are the number of neighbors ($K = [100, 200, \dots 500]$) 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 $[0, 5, \dots 25]$). 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. Since this distance and the selection of neighbors is sensitive to the scale of the features, we use min-max scaling for any continuous feature.

Finally, for the random forest estimates, we search over a parameter grid of the number of trees in the forest ("n_estimators"= $[100, 200, \dots 500]$) and the minimum impurity decrease for a split: ("min_impurity_decrease"= $[1, 2, \dots 10] \times 10^{-5}$).

B. ADDITIONAL TABLES AND FIGURES



Figure B.1: Nonperforming Loan Rate Decomposition

(a) Decomposition of Bank Nonperforming loans

(b) Nonperforming Loans by Loan Type

Notes: The left panel decomposes the bank nonperforming loan rate into delinquencies for construction and land development (blue), nonfarm nonresidential (red), and multifamily (green) loans. The right panel presents the equivalent nonperforming loan rates for each bank CRE loan category as well as the overall delinquency rate for CMBS loans. Nonperforming loans are loans that are 30 days or more past due or nonaccrual, plotted as a share of aggregate outstanding balances. CMBS calculations exclude defeased and REO loans.

Sources: Call Reports, Morningstar, and authors' calculations.



Figure B.2: Nonperforming Loan Rates by Bank Size





Notes: The figure plots CRE nonperforming loan rates for nonowner-occupied nonfarm nonresidential (NFNR) loans (left panels) and other CRE loans (right panels). Nonperforming loans are loans that are 30 days or more past due or nonaccrual, plotted as a share of aggregate outstanding balances. Other CRE includes multifamily, construction and land development, and owner-occupied NFNR loans. The top panels plot NPLs over time for banks above (blue) and below (green) \$100 billion in assets. The bottom panels plot 2023:Q4 NPLs against bank size. Blue dots report individual banks' NPLs, while the black line plots an estimate of the average for banks of that size (the kernel-weighted local mean). The scale on the y-axis expands after 20% to prevent outlier responses from obscuring variation within normal bounds. *Sources:* Call Reports and authors' calculations.





(a) NPL Decomposition by Property Type

(b) NPL Rates by Property Type

Notes: The left panel decomposes the bank nonperforming loan rate for income producing properties at large banks by property type, with the shaded region showing the contribution of a particular property type. The right panel presents the nonperforming loan rates for each property type. *Sources:* Y-14Q, and authors' calculations.





Notes: Decision tree estimated default probabilities. 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:* Y-14Q H.2 Schedule.



Figure B.5: Correlation Across Model Predictions

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





Notes: Figure plots the relationship between the weighted average fitted delinquency rate for the loans matched to a particular bank in the RCA/CoreLogic data and that bank's non-performing loan rate as of 2023:Q4. The red line shows the relationship based on a linear regression and the blue line based on local mean smoothing, both weighting by banks' volume of outstanding CRE loans. The dashed line gives the 45 degree line, or the relationship we would get with a perfect model. We average across the four delinquency models to produce the bank-specific fitted-deliquency rate.

Sources: Y-14Q H.2 Schedule, CoreLogic, MSCI RCA, and authors' calculations.

Data	Y-14	Morningstar	Y-14	Morningstar	RCA/Co	oreLogic
Sample	Large	CMBS	Large	CMBS	Large	Small
	Banks		Banks		Banks	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
Ν	42267	15532	42267	15532	133954	209810
Weighted			\checkmark	\checkmark	\checkmark	\checkmark

Table B.1: Average CRE loan characteristics Across Samples

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 3). 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/CoreLogic data. The measure of delinquency for the RCA/CoreLogic sample 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: Y-14Q H.2 Schedule, Morningstar, CoreLogic, MSCI RCA, Call Reports, and authors' calculations.

	$100 \times \text{Extension}_{i,23}$							
		Full Sample			Offices			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
CMBS	-5.37**	-6.40**	-6.21**	-6.26**	-6.26**	-9.74**	-9.41**	-9.42**
	(0.50)	(0.55)	(0.53)	(0.53)	(0.70)	(1.12)	(1.08)	(1.08)
Maturing	33.20**	32.27**	31.87**	31.60**	37.85**	36.59**	34.85**	34.70**
	(2.18)	(2.14)	(2.15)	(2.18)	(2.32)	(2.26)	(2.26)	(2.25)
Office	3.41**	1.43**	1.40**	1.33**				
	(0.47)	(0.40)	(0.38)	(0.37)				
LTV at Orig.		4.77**	4.60**	4.21**		9.43**	10.14**	9.81**
		(0.66)	(0.68)	(0.68)		(2.15)	(2.23)	(2.24)
ln(Value at Orig.)		2.10**	1.89**	1.80^{**}		2.58**	2.38**	2.35**
		(0.19)	(0.18)	(0.18)		(0.32)	(0.32)	(0.32)
Recourse		1.36**	1.47**	1.43**		-0.13	-0.07	0.00
		(0.26)	(0.27)	(0.26)		(1.05)	(1.02)	(1.03)
CBD			0.95**	0.91**			1.36	1.14
			(0.28)	(0.27)			(0.83)	(0.84)
Teleworkable Share			-1.94	-2.66			-3.64	-4.33
			(1.91)	(1.93)			(4.92)	(4.98)
Occupancy				-4.80**				-4.01*
				(1.03)				(1.92)
Debt Yield<.08				0.42^{+}				0.95
				(0.23)				(0.71)
R_a^2	0.231	0.246	0.241	0.243	0.256	0.275	0.262	0.264
Observations	59,179	59,095	54,326	53,616	7,849	7,829	7,189	7,064
Other Property Fixed Effects?	\checkmark	\checkmark	\checkmark	\checkmark				

Table B.2: Loan Modifications by Lender Type

Notes: This table presents estimates from the equation:

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

where Extension_{*i*,23} is an indicator for whether loan *i* received an extension (i.e., had its maturity date pushed out) in the last year. 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 the loan's at-origination LTV, whether the loan has recourse, and for 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 under 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:* Y-14Q H.2 Schedule, Morningstar, and authors' calculations.