Work-From-Home, Commercial Real Estate Risk and Credit Supply: Evidence from a Large Sample of Bank Loan Portfolios*

Elliot Anenberg† You Suk Kim‡ Erica Moszkowski§

April 1, 2025

Abstract

The rise in work-from-home (WFH) rates initiated by the pandemic has contributed to reduced commercial real estate (CRE) valuations and cash flows. This valuation decline increases credit risk in banks' CRE portfolios, which could impact the real economy if banks reduce lending (to CRE or other borrowers). Using a new data source, we calculate the exposure of thousands of individual banks to the WFH shock based on the geographic distribution of their outstanding pre-pandemic CRE loan portfolio. We find that banks with a one standard deviation increase in exposure decrease new CRE originations by 7.3 percent due to tighter credit supply. We document negative spillover effects across lending categories: banks with higher WFH exposure in their CRE portfolios also reduce credit supply for risky non-CRE loan types. Finally, we estimate that the average zipcode experienced a 25 pp decrease in the growth rate of total new originations due to the bank lending response to the WFH shock. These results suggest some but incomplete substitution of credit demand from exposed banks to nonbanks.

*The analysis and conclusions contained in this presentation are those of the authors and do not necessarily reflect the views of the Board of Governors of the Federal Reserve System, its members, or its staff.

†Federal Reserve Board of Governors. Email: elliot.anenberg@frb.gov.

‡Federal Reserve Board of Governors. Email: you.kim@frb.gov

§University of British Columbia. Email: erica.moszkowski@sauder.ubc.ca

1 Introduction

A seemingly long-lasting economic consequence of the COVID-19 pandemic is a stark increase in work-from-home (WFH) rates (Barrero et al., 2023). The rise in WFH has the potential to bring important welfare benefits in the long run, for example through reduced commuting, improved housing affordability, or higher productivity.[1] A salient shorter-run consequence of the shift to WFH, however, has been a large drop in commercial real estate (CRE) cash flows and valuations (Gupta et al., 2022), with potentially important adverse consequences for the financial system and the real economy (Jiang et al., 2023).[2]

A key channel through which these valuation declines could have broader impacts on the economy is through banks' exposure to loans backed by distressed CRE properties. CRE makes up about a quarter of total assets at the average bank, and as of the end of 2023, banks held about $2.8 trillion of the $5.8 trillion in CRE loans (including multifamily) outstanding. An increase in expected delinquencies and defaults in a key part of the bank's risk-asset portfolio may incentivize the bank to reduce risk by tightening credit supply. If reduced credit supply from banks is not offset by nonbanks, then the decrease in bank credit supply for CRE would likely have important negative effects on the real economy, for example through lower construction of CRE properties (Peek and Rosengren (2000)) or lower corporate investment (Chaney et al. (2012)). Furthermore, if banks reduce other (non-CRE) types of lending to maintain their overall capitalization ratios, real effects could be more widespread. These credit supply effects are likely to play out over the next several years, as most measures of work-from-home rates have stabilized well above pre-pandemic levels[3]

[1]See, for example, Delventhal and Parkhomenko (2023), Howard et al. (2023), Delventhal et al. (2022), Davis et al. (2024).

[2]CRE prices fell 21 percent peak to trough according to the CoStar repeat sales index. While the price drop has been most severe for office properties, other commercial property types have also been adversely affected by WFH and have experienced price declines due to lower foot-traffic from office workers, reduced business travel, or shifting residential housing preferences (e.g. DeFusco et al. (2023), Mondragon and Wieland (2022)).

[3]For example, Barrero et al. (2021)'s Survey of Working Arrangements and Attitudes finds that about 26% of paid work days in the U.S. in February 2025 were work-from-home days, and this number has remained stable since January 2022.

and banks' CRE portfolios are made up predominantly of loans originated at pre-pandemic valuations.

In this paper, we empirically study the effects of the WFH shock induced by the pandemic on credit availability through this bank exposure channel. Our first step is to compute a bank's "exposure" to the WFH shock for a large sample of banks. To do so, we leverage public records data from CoreLogic, which allows us to infer the detailed location of the properties collateralizing CRE loans on each bank's pre-pandemic balance sheet. We use the share of pre-pandemic county employment that could possibly be done remotely, as measured by Dingel and Neiman (2020), as a local area's exposure to the WFH shock. A bank's exposure is a weighted average of the share of jobs that can be done remotely in each county in which the bank has outstanding CRE loans. The weight for each county equals the share of the bank's outstanding CRE portfolio backed by properties in that county. For example, a bank with a high share of its pre-pandemic CRE loans collateralized by properties in San Francisco county, a well-known example of an area with an industry mix that allowed for high WFH rates once the pandemic hit, would be highly exposed under our definition.

We estimate the effect of bank exposure on bank CRE credit supply using an identification strategy similar to Khwaja and Mian (2008). Essentially, we compare levels of CRE mortgage originations among banks lending in the same market (county-year) with varying levels of exposure. Importantly, market-level credit demand shocks, which could be correlated with bank exposure, are partialed out by fixed effects. Our regression also controls for various bank characteristics such as the bank size, the share of CRE loans in total assets, and the share of CRE loans collateralized by owner-occupied properties.

We find that banks with greater exposure to the WFH shock sharply reduced CRE lending after the pandemic. In the years leading up to the pandemic, we show no evidence of any pre-trend in the relationship between Exposure and lending. Starting in 2020, the first year of the pandemic, the relationship turns negative and stays negative throughout the post-pandemic period. On average across all post-pandemic years, a one standard deviation increase in Exposure is associated with a 7.3 percent decrease in the number of originations. We find similar but slightly noisier effects on origination volumes.

Our proposed mechanism for these results is that more-exposed banks experience larger (expected) losses in their CRE loan portfolio, causing them to reduce credit risk in their loan portfolio by tightening credit supply for new originations of credit, especially loans against riskier CRE properties. Since CRE mortgages are very illiquid assets, one of the easiest ways for banks to adjust their CRE portfolio risk is through a change in origination behavior. Indeed, we find that more-exposed banks indeed have a higher level of CRE loan delinquencies, using bank regulatory filings (Call Reports) merged to our CoreLogic data. We also present evidence that more-exposed banks reduced their lending for CRE properties in counties with larger WFH shocks, which likely experienced worse declines in CRE fundamentals and are plausibly riskier counties for banks to lend in.

The incentive for de-risking is likely to be be larger if CRE loan losses can lead to a larger negative shock to bank capital. A large negative capital shock may shift banks closer to a constraint for raising new external funds as in the model of Froot and Stein (1998), and banks may especially need to de-risk to comply with regulatory requirements. Consistent with this hypothesis, we find that banks with a larger CRE loan portfolio relative to their Tier 1 capital reduce their lending by more given the same level of exposure to the WFH shock. We also find the very largest banks reduce their lending by less, even after controlling for capitalization and the number of markets they lend in. This is consistent with the long literature on bank risk aversion: if larger banks are less risk averse than small banks, either because they are too big to fail (Naqvi and Pungaliya (2023)) or because they can more easily cover realized losses due to a low cost of raising deposits ((Kashyap and Stein, 1995)), they may not pull back on new originations as much. Finally, we show that derisking is not confined to the CRE portfolio. More exposed banks tighten credit supply for non-CRE, risky loan types including commercial and industrial loans.

In the last part of our paper, we estimate whether the pullback in CRE credit supply by exposed banks led to a decrease in credit supply at the aggregate, market level. It is ex ante unclear whether bank exposure to the WFH shock led to a decline in the aggregate credit supply. Some CRE loan borrowers might have substituted from exposed banks to nonbanks or less-exposed banks, especially as credit from capital markets was widely available over most of our sample period. Our estimates show some evidence of substitution from exposed banks to nonbanks, but we nevertheless find substantial aggregate credit supply effects. We estimate that the average zipcode experienced a 25 percentage point decrease in the growth rate of total new originations (bank + nonbank) due to the bank lending response to the WFH shock. This represents the combined effects of a 56 percentage point decrease in the growth rate of originations by banks and a 15 percentage point increase in the growth rate of nonbank originations (again in the average zipcode).

Our analysis relies on representative measurement of the CRE lending activities of thousands of banks, including small and regional banks, at a granular geographic level. Traditional CRE datasets used in the literature are limited for our purposes because they only have information for a handful of large banks or only cover bank originations of large loans.[4] Rather than relying on these datasets, we construct a new dataset based on CoreLogic's Real Estate Public Records data. Our constructed data (i) provide a large enough sample of banks for us to use a Khwaja-Mian estimator to credibly separate credit supply effects from credit demand effects and (ii) are more representative so that we can compute more accurate measures of bank exposure for our sample of banks.

Related literature. Our paper contributes to the growing literature on the implications of work-from-home for CRE valuations and the broader economy. Rosenthal et al. (2022) document how the pandemic changed the value firms place on access to city centers. Gupta et al. (2022) estimate how much the WFH shock leads to a decline in the office property

[4] Regulatory datasets, as used in (Glancy et al., 2022a, 2023), contain only the largest banks. The Real Capital Analytics data used in (Ghent and Valkanov, 2016; Büchler et al., 2024) do not cover loans below $ 2.5 million. Given a lack of loan-level data on bank balance sheet CRE loans, Jiang et al. (2023) turn to CMBS data, which are richer and more comprehensive, as a proxy to gauge potential distress in bank CRE loans.

value in New York City. Jiang et al. (2023) estimate the effects of the CRE property value decline and the monetary policy tightening on CRE loan value and simulate the effects of the CRE distress on bank solvency run risk. Glancy and Wang (2024) examine the adverse effects of lease expiration on office property fundamentals. Glancy and Kurtzman (2024) study the determinants of post-pandemic CRE loan delinquencies, finding that small banks' comparatively modest delinquency rates mostly reflect observable portfolio characteristics (e.g. their low holdings of large-sized office loans). Our paper contributes to this literature by providing empirical estimates of the effects of the WFH shock on the supply of CRE loans as well as other lending products.

Our paper is also related to a literature showing that shocks to bank net worth influence their choice of new risk-assets. This effect is often called the "bank lending channel" (Khwaja and Mian (2008), Bidder et al. (2021), Chodorow-Reich (2014), Ramcharan et al. (2016), Ivashina and Scharfstein (2010)). We contribute to this literature by focusing on a shock to bank net worth coming through the bank's CRE portfolio. Much of the existing, recent literature on the bank lending channel focuses on corporate lending. In addition, like Greenstone et al. (2020), we focus on shocks to bank balance sheets that occurred amid a strong macroeconomy whereas much of the existing literature focuses on times of crisis. Even during good economic times, we find evidence of a strong bank lending channel that has important consequences for aggregate credit supply.

2 Data and Context

2.1 CoreLogic Real Estate Data

We measure bank exposure to commercial real estate loans at a geographically disaggregated level using public records data from CoreLogic. These data cover the universe of property transactions, including information about the mortgage if there is one, in covered counties. Data are sourced from public records filed at the county deeds office. Mortgage refinancings are also recorded in this dataset, since they expunge an existing lien on a property and replace it with a new one.

Though it has been used extensively to study residential real estate markets, few papers use CoreLogic to study commercial real estate. Babalievsky et al. (2023) use CoreLogic's tax data; we instead use CoreLogic's Owner Transfer and Mortgage Basic files. By contrast, most papers in the in the commercial real estate literature rely on loan-level data from bank Y-14 filings (Glancy et al., 2022b), Real Capital Analytics (Büchler et al., 2024; Glancy and Kurtzman, 2024), or Trepp (Jiang et al., 2023). These datasets have higher quality information on property types and lender identities than CoreLogic, but much more limited coverage. Y-14 filings are only required for banks with assets totaling more than $100 billion, so it is not a useful data source for studying small and regional banks. Real Capital Analytics only covers properties valued at $2.5 million or higher, which are more likely to be originated by large banks. Trepp covers only loans which have been securitized into CMBS.

The main disadvantage of the CoreLogic dataset is that it does not track loan performance. As a result, we are unable to observe whether any particular loan becomes delinquent or is defaulted upon. To measure the performance of a bank's overall CRE portfolio, we turn to bank regulatory data.

2.2 Bank Regulatory Data

We also use quarterly Call Reports and FR Y9-C filings to obtain bank-level information that is not available in CoreLogic. Both data sets are drawn from regulatory filings from banks or bank holding companies (BHCs). Call Reports are regulatory filings which provide information about a bank subsidiary of a BHC (for example, Wells Fargo Bank), and FR Y9-C provides information about a BHC (for example, Wells Fargo & Company). Both data sets report detailed information about bank balance sheets and incomes at a quarterly frequency. These data do not allow for geographic disaggregation, but provide a comprehensive picture of a bank's total size, capitalization, and the size and performance of its CRE portfolio, including including amounts of loans in delinquencies, amounts charged offs, and amounts in allowances in anticipation of future losses.

2.3 Match between CoreLogic and bank regulatory filings

Although CoreLogic provides lender information for each loan, it has a few limitations. First, it is difficult to link loans originated by the same lender within CoreLogic. Internal lender identifiers available in CoreLogic seem to be only minimally cleaned. Thus, slightly different names of the same bank – for example, "Wells Fargo Bank" and "Wells Fargo BK NA" – often have different internal identifiers. Second, it is also difficult to match the data to other lender-level data sets like Call Reports because CoreLogic does not provide additional information about lenders such as their tax IDs, which can be used to match with another data set.

Therefore, to aggregate loans to the bank level, we match CoreLogic lender names to RSSD IDs, which is a unique identifier assigned to financial institutions by the Federal Reserve. Using RSSD IDs, we can combine CoreLogic with Call Reports and FR Y9-C.

Details about the data match are described in Appendix A. Overall, we match about 91% of bank CRE loans in CoreLogic to unique RSSD IDs. In terms of loan amount, about 99% of bank CRE loans in CoreLogic are matched to RSSD IDs, suggesting our match rate is higher for larger banks, which are likely to originate large loans.

2.4 Summary Statistics of Loan Characteristics in CoreLogic

Table 1 presents summary statistics of characteristics of CRE loans originated by banks from 2016 to 2023 in CoreLogic. As discussed in detail later, we use CRE loan information in CoreLogic to calculate banks' exposures to WFH shocks through locations of properties backing CRE loans on bank balance sheets. Unfortunately, CoreLogic does not allow us to determine which loans were securitized into CMBS, although some banks are active in CMBS securitization. Thus, the sample excludes loans originated by 27 banks that are active in

CMBS loan originations using the RCA data.

The average loan original balance is about 2 million, but the median is much smaller at around 550,000 due to a small fraction of very large loans. The information on loan maturity is non-missing for about 40% of loans in the sample, and the average loan term is almost 15 years, while the median is 7 years. Some loans have a very long maturity of up to 30 years. The average mortgage rate in our sample is 5.3%, but the information on mortgage rates is missing for a majority of the sample.

In our sample, there is no particular concentration in a single property type. Multifamily and retail properties have the highest shares of around 17%. Industrial and office properties also have non-negligible shares of 13% and 10%. Properties with multiple types ("mixed used"), such as office or apartment buildings with store fronts, account for 10%. Other property types include specialty properties like hospitals and other medical buildings and have the share of 15%. For 16% of CRE loans with unspecified types, CoreLogic does not provide further information about their property types.

Lastly, 30% and 25% of originations are purchase and refinance loans, respectively. Junior-lien loans have a large share of 38%, and 8% of loans have unknown loan types.

Table 1: Loan-Level Summary Statistics

	Mean	SD	p5	p50	p95	Count
Original Loan Balance (millions)	2.04	4.67	0.06	0.55	8.87	618,837
Mortgage Term (years)	13.50	9.36	3.00	10.00	30.00	250,376
Mortgage Rate (%)	4.77	1.61	2.62	4.60	7.50	29,570
Property Type:						
Office	0.10	0.30	0.00	0.00	1.00	641,290
Multifamily	0.17	0.37	0.00	0.00	1.00	641,290
Industrial	0.13	0.34	0.00	0.00	1.00	641,290
Retail	0.17	0.37	0.00	0.00	1.00	641,290
Hotel	0.02	0.13	0.00	0.00	0.00	641,290
Mixed Use	0.10	0.31	0.00	0.00	1.00	641,290
Other Types	0.15	0.36	0.00	0.00	1.00	641,290
Types Unspecified	0.16	0.37	0.00	0.00	1.00	641,290
Loan Type:						
Purchase	0.30	0.46	0.00	0.00	1.00	641,290
Refi	0.25	0.43	0.00	0.00	1.00	641,290
Junior-Lien	0.38	0.48	0.00	0.00	1.00	641,290
Unknown	0.08	0.27	0.00	0.00	1.00	641,290
Number of Observations	641,290					

Note: Characteristics of CRE loans originated by banks in the period from 2016 to 2023. The sample exclude loans originated by banks that are active in CMBS securitization. Source: CoreLogic.

2.5 Computing loans outstanding in 2019Q4

CoreLogic only provides information on loans at the time of origination, so to determine bank CRE exposure just before the onset of COVID-19 we need to infer how much principal remains on bank balance sheets at the end of 2019. This requires converting each loan's origination record into a quarterly panel tracking the amount of principal outstanding over time.

Our first step is to filter out loans which are likely no longer on banks' books at the end of 2019. This is difficult for two main reasons. First, refinancing prior to maturity seems quite common, so relying on the loan's maturity date to filter out old loans does not remove all loans which are no longer on banks' books. Second, many commercial properties have both first and second lien mortgages which we can observe in the data. This makes it hard to determine whether a new loan on the same property is a refinancing of a preexisting mortgage or adds an additional lien. We keep only loans which have a maturity date after 2019Q4 and which are the last mortgage originated with its own seniority level. For example, if a property has a senior and junior lien outstanding in 2017Q4 and then we observe a new senior mortgage in 2018Q1, we drop the original senior mortgage.

Next, we compute the amortization schedule of each loan to determine the amount of principal outstanding in 2019Q4. For fixed-rate mortgages, we apply the standard amortization formula to compute quarterly mortgage payments, and then break out the interest and principal components of that payment. Since the mortgage rate is often missing in CoreLogic, we impute the average mortgage rate for other fixed-rate mortgages originated in the same quarter. Moreover, we replace missing maturity information with the most common maturity among loans with the same property type.[5] All the data fields required to calculate exact quarterly payments and amortization schedules for adjustable-rate mortgages (the length of the locked-rate period, the frequency of rate adjustments, minimum and maximum adjustment intervals, and minimum and maximum rates) are not usually available. Instead, we approximate the amortization schedule of adjustable-rate loans with a fixed-rate amortization schedule based on the starting mortgage rate.

Finally, we account for lender exit, whether through bankruptcy or merger. If the lender exits at any point between loan origination and 2019Q4, we reassign the loan to the acquiring firm. If there are multiple acquiring firms (as sometimes occurs after a bank failure), we

[5]For most property types, this is a fairly long loan term of 10-30 years. However, our measure of outstanding balances is not terribly affected by imputing a loan term which is too long for two main reasons. First, most commercial mortgages have amortization periods which are longer than the term of the loan itself. Though data on amortization periods is not always available, industry professionals usually cite a standard amortization period of 30 years. This is exactly the loan term we impute for most property types. Second, quarter-to-quarter changes in the dollar value of loans outstanding are second-order in magnitude compared to adding and removing loans from the balance sheet in their entirety. Most properties either refinance or sell well within the imputed term of the loan, and we observe both types of transactions in CoreLogic.

Figure 1: Benchmarking CoreLogic against the Call Report

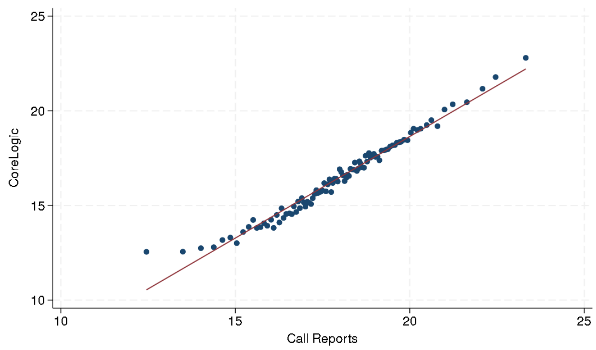

(a) Log of CRE loans outstanding as of 2019Q4

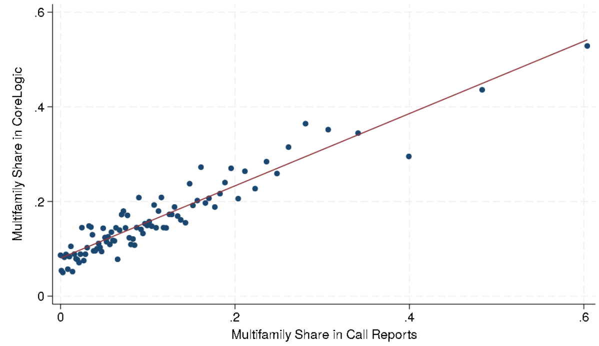

(b) Ratio of multifamily loans outstanding to CRE loans outstanding as of 2019Q4

Note: Include both MF and non-MF loans. Excludes CMBS banks.

assign each loan to each acquirer proportionally to the number of acquirers.

This process results in a panel of loans outstanding as of 2019Q4 for each bank. We can compare some aspects of a bank's CRE loan portfolio to their counterparts in the Call Report to gauge whether our process to construct the loan portfolio with CoreLogic is reasonable. Although the Call Report provides a very limited set of characteristics of a bank-level CRE loan portfolio, we can still observe the size of the CRE loan portfolio and a breakdown of loans for multifamily properties and other CRE loans. The binned scatter plot in Figure 1(a) shows that the CRE loan amount outstanding as of 2019Q4 constructed with CoreLogic is highly correlated with its counterpart in the Call Report. Although the figure also shows that the CoreLogic-based portfolio tends to underestimate the CRE portfolio size in the Call Report, banks with varying sizes of CRE loan portfolios appear to be affected similarly, except for banks with the smallest portfolios. This suggests that the underestimation of the portfolio size is unlikely to disproportionately impact banks with different levels of WFH exposure, [which we confirm later in Section 3.] Moreover, the binned scatter plot in Figure 1(b) shows that the share of multifamily loans in the CRE portfolio is highly correlated with CoreLogic and the Call Report. These two figures suggest that a bank's CRE loan portfolio based on CoreLogic has characteristics comparable to information in the Call Report.

2.6 Comparing CoreLogic and RCA

One of our contributions is to shed light on CRE lending activities by smaller banks using the CoreLogic data, in contrast to other CRE data sets covering only a certain segment of the CRE lending market. CoreLogic is the natural data source for our analysis, since we are interested in measuring small bank exposure to commercial real estate in different geographic areas. This requires data with sufficient coverage of both small banks and small loans, which other datasets lack.

In particular, the CRE loan origination data set from Real Capital Analytics (RCA) has been used by some previous papers, such as Büchler et al. (2024), Ghent and Valkanov (2016), and Glancy et al. (2022a). However, the RCA dataset only covers CRE properties with valuations of at least 2.5 million dollars. Figure 2 shows that RCA covers roughly half of CRE lending by banks in CoreLogic. For this calculation, we exclude banks that are heavily involved in CMBS loan originations (those those with at least 10% of their originations for CMBS) to focus on the difference in data coverage on loans funded by bank balance sheets. Because the geographical coverage of CoreLogic has expanded over time, we also restricted the RCA sample to county$\times$years that are also covered by CoreLogic.

A comparison between the CoreLogic and the Real Capital Analytics datasets sheds light on how much smaller banks' activities are missing in non-comprehensive data sets. In recent years, the aggregate lending amount against CRE properties by banks in RCA is roughly a little more than half of that in CoreLogic. Although smaller-sized CRE loans (especially below $2 million) account for the vast majority of CRE lending in terms of number of loans in CoreLogic, many of these loans are missing in RCA because it covers relatively large transactions.

Figure 2(a) shows that aggregate bank lending volume in RCA is about 55% to 90% of that in CoreLogic depending on the years. The average ratio of the RCA volume to the CoreLogic volume between 2015 and 2023 is about 71%. The ratio is even smaller in terms of loan count (Figure 2(b)), suggesting that RCA's the minimum transaction size threshold limits its coverage of smaller loans.

Figure 2: Aggregate bank CRE loan originations

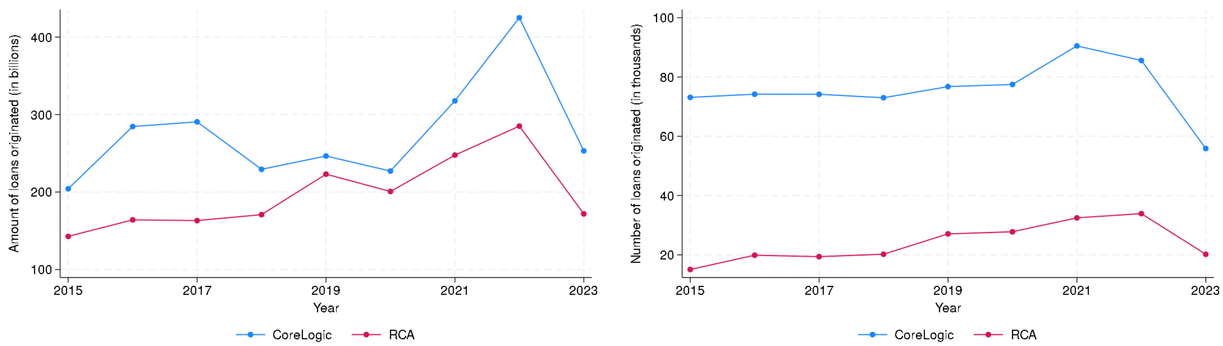

(a) Aggregate amount of loans originated (b) Aggregate number of loans originated

Note: Bank loans only. Dropped CMBS lenders from both data sets. The RCA sample was restricted to to county×years that are also covered by CoreLogic. Authors' calculations based on CoreLogic and RCA.

Figure 3 shows more directly that many small loans, especially below \$2 million, are missing in RCA, but not in CoreLogic. These relatively small loans are much more likely to be originated by smaller banks. Banks with total assets below \$1 billion account for 32% of loans below \$2 million and 18% of loans above \$2 million in 2022. In contrast, large banks with more than \$10 billion in their assets account for 31% of loans below \$2 million and 45% of loans above \$2 million in 2022.

Figure 3: Distributions of bank loan size in 2022 in CoreLogic and RCA

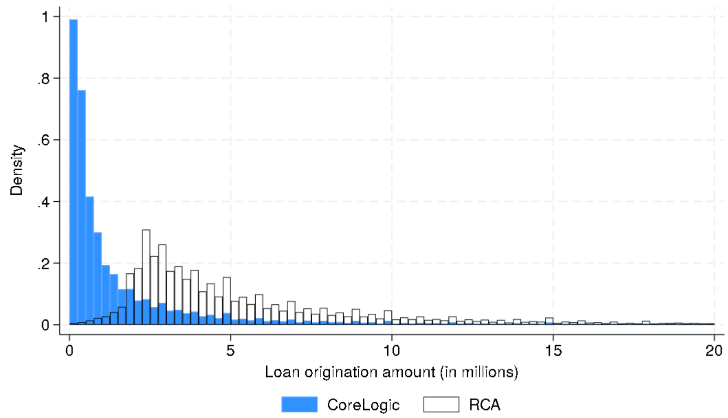

Note: Bank loans only. Dropped CMBS lenders from both data sets. The RCA sample was restricted to to county×years that are also covered by CoreLogic. Truncated loans larger than 20 million dollars, which account for about 1% of the sample. Authors' calculations based on CoreLogic and RCA.

3 Empirical strategy and main results

We examine the causal effects of the WFH shock on bank CRE credit supply. We compare banks with different exposures to this shock, using regional variation in locations of properties backing pre-shock bank CRE loans and variation in how much different regions are subject to the WFH shock.

3.1 Measuring bank exposure to WFH shock

We measure bank i's exposure as a weighted average of the county-level share of jobs that can be done remotely, weighted by the bank's pre-shock outstanding CRE loan share in each county (c):

$$Exposure_i = \sum_c \frac{\text{CRE loans outstanding}_{ic}}{\text{total CRE loans outstanding}_i} \times WFHshare_c \tag{1}$$

Intuitively, our exposure metric is a Bartik shock, where the share of jobs that can be done remotely is the "shock" component, and CRE loans are the weights that determine how exposed each bank is to the shock. We measure CRE loans outstanding as of 2019Q4 based on CoreLogic data.

$WFHshare_c$ is the share of employment in county c that can be done remotely, defined as:

$$WFHshare_c = \frac{\sum_k emp_{kc} \times WFH_k}{\sum_k emp_{kc}}, \tag{2}$$

where emp_{kc} is 2019Q4 employment in industry k in place-of-work county c from the Quarterly Census of Employment and Wages (QCEW).[6] WFH_k is the share of employment in industry k (defined at the 3-digit NAICs code level) that is teleworkable, as measured in Dingel and Neiman (2020). Dingel and Neiman (2020) use surveys (conducted before the pandemic) describing the typical experience of US workers to classify each occupation as able or unable to be done entirely from home.[7] For example, if a majority of respondents to the survey in a given occupation report that they work outdoors every day, then the occupation is coded as cannot be done from home. The industry with the lowest and highest level of WFH_k is "food services and drinking place" (0.0175) and "Securities, Commodity Contracts, and Other Financial Investments and Related Activities" (0.95), respectively.

Figure 4(a) shows substantial variation in $WFHshare_c$ across counties, ranging from below 10% to more than 50%. Figure 4(b) shows that variation in our measure of bank exposure to the WFH shock $Exposure_i$ in equation 1 is also large across banks. Because $Exposure_i$ is a weighted average of $WFHshare_c$, these two figures suggest that there is substantial heterogeneity in regions in which banks lend in terms of the share of employments done remotely.

We next explore correlations between Exposure and bank characteristics. Table 2 presents summary statistics for several bank characteristics measured from the Call Reports for the

[6]Note that a single firm can have employees working across multiple locations and/or industries. For example, Amazon employees working in Seattle may be in the tech industry but in most other work locations could be in the warehousing industry.

[7]Dingel-Neiman also provide a measure of WFH_k that does not use the surveys, but is based on the two authors' manual inspection of occupation descriptions. The two measures are highly correlated and our results are essentially unchanged using this alternative measure.

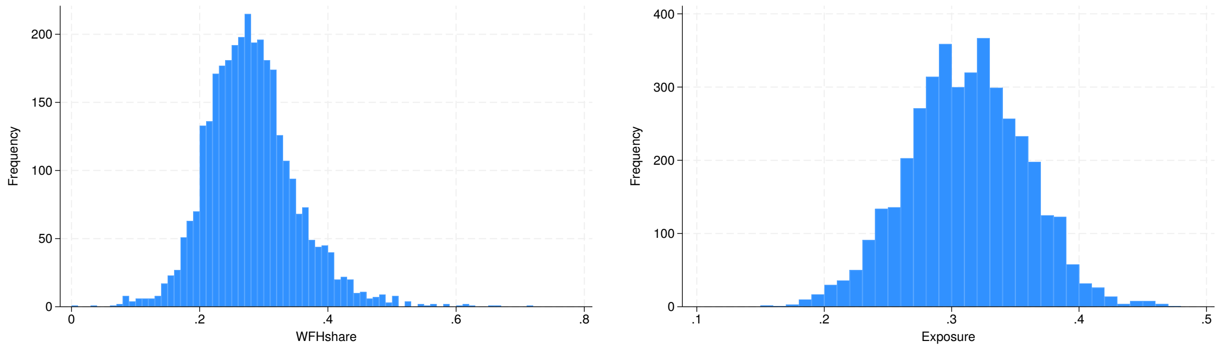

(a) Variation in $WFHshare_c$ across counties (b) Variation in $Exposure_i$ across banks

Note: Authors' calculations based on the data from QCEW and CoreLogic.

nearly 2000 banks in our main estimation sample. We use these bank characteristics as control variables in our main regressions described below. Table 3 presents a regression of the z-score of Exposure (i.e. standardized so that it has zero mean and one standard deviation) on those bank characteristics. The regression shows that exposure is negatively correlated with bank size. Banks with less than \$ 1 billion in assets have a 0.82 lower standardized exposure (i.e. 82 percent of one standard deviation of unstandardized exposure) compared to the very largest banks, those with over \$ 100 billion in assets. More exposed banks also tend to have higher capital ratios; a lower share of owner-occupied CRE properties, which were less adversely affected by the WFH shock; and a low share of debt securities relative to total assets, all else equal.

Table 2: Summary Statistics

	mean	sd	p5	p50	p95	count
Exposure	0.32	0.04	0.25	0.32	0.38	1944
CRE loans outstanding (billions of dollars)	470.82	3438.47	1.82	39.70	1244.19	1944
Total assets (billions of dollars)	3.07	20.05	0.09	0.45	7.95	1944
CRE assets share of Tier 1 capital	1.64	20.17	0.09	0.80	3.26	1944
Owner-occupied share of CRE assets	0.49	0.21	0.12	0.48	0.87	1944
Tier 1 capital ratio	0.16	0.06	0.11	0.14	0.27	1944
Debt securities share of total assets	0.16	0.12	0.01	0.14	0.39	1944
Number of counties with CRE loans	25.22	54.77	2.00	11.00	95.00	1944

Note: Bank characteristics measured as of 2019Q4 using Call Reports data.

Table 3: Exposure and Bank Characteristics

	(1)	
	Exposure	
	Coef.	SE
10 - 100 billion dollars in assets	-0.096	(0.355)
1 - 10 billions dollars in assets	-0.374	(0.385)
$<$ 1 billion dollars in assets	-0.818**	(0.393)
CRE loans outstanding (billions of dollars)	0.000	(0.000)
CRE assets share of Tier 1 capital	0.000	(0.001)
Owner-occupied share of CRE assets	-0.890***	(0.102)
Tier 1 capital ratio	1.264***	(0.354)
Number of counties with CRE loans	0.000	(0.001)
Debt securities share of total assets	-1.691***	(0.190)
R-squared	0.146	
N	1944	

Note: Omitted group is banks with over 100 billion dollars in assets. Bank characteristics measured as of 2019Q4. Exposure is standardized to have zero mean and standard deviation equal to one.

3.2 Research design

In our main analysis, we estimate the effect of a bank's exposure to the WFH shock on its county-year-level CRE purchase mortgage originations. We use a Poisson regression instead of linear regressions of log of origination volume because there are many bank-county-year observations with zero CRE loan originations, and the log of origination volume would be undefined for these observations.

We assume that the count of loan originations at bank i, in market j (county), and year t follows a Poisson distribution with conditional mean

$$\lambda_{i,c,t} = \exp(\beta_t \text{Exposure}_i + X_i\gamma + \delta_{c,t}), \tag{3}$$

where β_t are year dummies interacted with the z-score of Exposure, and are the coefficients of interest. The sample period is from 2016 to 2023, and we normalize β_{2019}, the coefficient associated with the year just before the pandemic starts, to zero. These coefficients capture a percentage change in the number of CRE loan originations in response to a one standard deviation increase in Exposure. Standard errors are clustered by bank.[8]

The bank controls X_i are listed in Table 2 and are measured as of 2019 Q4, right before the start of the pandemic. For the bank controls CRE loans outstanding and total assets, we include cubic polynomials of these variables as controls. We use the same set of bank controls in each specification described in the remainder of the paper.

$\delta_{c,t}$ are county-by-year fixed effects, which absorb variation in CRE lending due to location-specific demand or supply factors that are common across banks. For example, heterogeneous trends in CRE demand across counties, possibly due to different degrees of WFH shocks, do not affect the identification of β_t. Instead, the variation pinning down β_t is within-county-year variation in *Exposure*, which comes from variation in the geographic composition of a bank's CRE loan portfolio as well as variation in size of the WFH shock

[8]In the Appendix, we show results for the count of refinance originations and the dollar volume of purchase originations.

across counties. We interpret the β_t's as reflecting credit supply responses to Exposure. We provide additional evidence that our estimates of β_t reflect credit supply and not credit demand in Section 5. Our identification strategy is similar to the one developed in Khwaja and Mian (2008), although they compared banks lending to the same firm, not to the same local market.

3.3 Effects of WFH exposure on bank CRE lending

We find that banks with greater exposures to WFH shocks sharply reduced their CRE lending after the pandemic. Figure 5 presents the estimates of β_t in equation 6 for the number of new bank CRE purchase loan originations. In each of the four years leading up to the pandemic (2016-2019), there is no evidence of any pre-trend in the relationship between Exposure and lending, conditional on our controls. Starting in 2020, the first year of the pandemic, the relationship turns negative and stays negative throughout the post-period. On average across all post-pandemic years, a one standard deviation increase in Exposure is associated with a 7.3 percent decrease in originations (column 1 of Table 5). The point estimates for the post-period vary slightly from year to year, with stronger effects in the high interest rate environment of 2023.

How much did the average bank restrict credit supply in response to the WFH shock? We calculate this as $\frac{0.073\times(0.32-WFH^*)}{0.042} = 47\%$, where 0.073 is the point estimate from column 1 of Table 5, 0.32 and 0.042 are the average and standard deviation of unstandardized bank exposure to WFH from row one of Table 2, and WFH^* is the average pre-pandemic WFH rate. We set $WFH^* = 0.05$ based on calculations from the 2019 American Community Survey (ACS), which is also consistent with evidence from the American Time-Use Survey (Barrero et al. (2021)).[9]

[9] Using the 5-year 2023 ACS, we calculate the actual WFH share using the share of workers aged 16 and over with at least \$10,000 in earnings that reported “worked at home" as their usual mode of transportation to work. We drop respondents who live in a county that is not in our main estimation sample. This measure of work-from-home is used in Mondragon and Wieland (2022) and discussed in detail in Buckman et al. (2025).

Figure 5: Effects of Exposure on bank CRE purchase loan originations

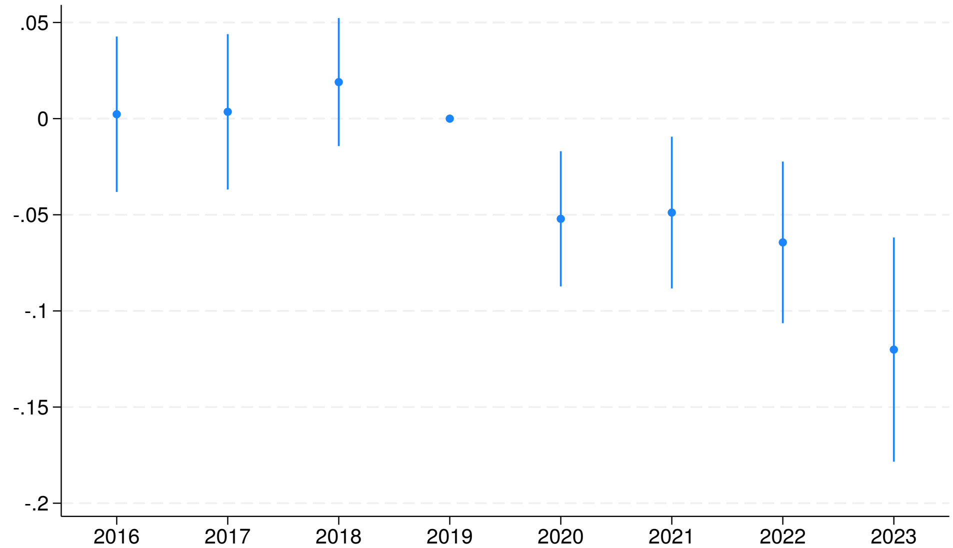

Note: This figure presents the estimates and the standard errors of β_t in Equation 6. Exposure is standardized to have mean of zero and standard deviation of one. The standard errors are clustered by bank.

It is difficult to put our results in context with existing literature, since the nature of the work-from-home shock is quite different than other shocks which have been studied. However, our estimates are still comparable to previous estimates of how banks responded to negative shocks to their loan portfolio. Bidder et al. (2021) find that a 1 standard deviation increase in the share of loans accounted for by oil and gas firms before the 2014 oil price shock was associated with a 4 percent decrease in the size of commercial loans. In the context of the Japanese banking crisis of the early 1990s, Peek and Rosengren (2000) find that a 1 standard deviation increase in the nonperforming loan ratio at Japanese parent company was associated with 1.2 percent decline in the growth rate of CRE loans outstanding.

4 Mechanism

Our interpretation of the results in Figure 5 is that more-exposed banks restrict credit supply by more than less-exposed banks. The mechanism we propose is that more-exposed banks experienced a distress in their CRE loan portfolio via an increase in current or expected future losses of their CRE loans, which in turn may incentivize banks to reduce credit risk in their loan portfolio by pulling back on new CRE lending, in particular against riskier CRE properties. CRE loans on bank balance sheets are illiquid and long-term loans, so origination is an important margin through which banks can reduce risk in the portfolio over time. The incentive for de-risking is likely to be larger if CRE loan losses can lead to a larger negative shock to bank capital. A large negative capital shock may shift banks closer to a constraint for raising new external funds as in the model of Froot and Stein (1998), and banks may especially need to de-risk to comply with regulatory requirements.

In this section, we perform several analyses to test our proposed mechanism. First, we explore the relationship between bank exposure and several direct measures of distress in banks' CRE portfolios. Second, we look at several dimensions of heterogeneity which could amplify or dampen the degree to which bank exposure to the WFH shock affects their net worth. Third, we test for spillover effects of bank exposure to other, non-CRE types of lending.

4.1 Exposure and loan distress

The WFH shock is only likely to cause distress in bank CRE loan portfolios if CRE valuation declines are large enough to bring property values below outstanding loan balances. Several existing papers find evidence that lenders are exposed to losses as a result of property valuation declines since 2019. Glancy and Kurtzman (2024) find that CRE nonperforming rates have increased since the pandemic, especially for CMBS and loans held by large banks. Using data from the CMBS market, Jiang et al. (2023) estimate that 14% of all CRE loans and 44% of office loans have property values below outstanding debt balances, exposing lenders to losses. In this section, using our Exposure metric, we test whether banks whose loan portfolios are more geographically exposed to the WFH shock experienced larger distress in their loan portfolios.

A natural measure of distress is the CRE delinquency or nonperforming loan rate. Our CoreLogic data, however, do not report any information on loan performance. Instead, we rely on the Call Report regulatory filings, which are available at the bank level (i.e. without any of the geographic detail available in the CoreLogic data). We calculate nonperforming CRE loans for each bank as the dollar volume of CRE loans at least 90 days past due or in nonaccrual in 2023Q4. The Call Report data also include information about loan loss provisions and chargeoffs, which can also be informative about bank distress. Loss provisions are a stock of funds set aside by banks to cover expected future credit losses. When losses are realized (for example, the property is foreclosed upon and sold), the losses are recorded as chargeoffs, and the stock of provisions decreases by the amount of chargeoffs.[10] We measure loss provisions in 2023Q4 and cumulate all chargeoffs between 2020Q1 and 2023Q4. Because the Call Report data report CRE-specific loss provisions and chargeoffs only for banks with assets of at least \$1 billion, we use loss provisions and chargeoffs are reported as totals for the bank across all types of loan categories (CRE and others). Thus, our measures of loss provisions and chargeoffs are noisy measures of actual loss provisions and chargeoffs for CRE loans. We divide all of our measures of distress by the total dollar volume of CRE loans on the bank balance sheet in 2023Q4.[11]

We regress our measures of distress on bank exposure using the following bank-level regression:

$$Y_i = \beta \text{Exposure}_i + \beta_X X_i + \epsilon_i, \tag{4}$$

where Y_i is one of our distress measures for bank i. We also include all of the bank controls discussed above, X_i.

Results can be found in table 4. A one standard deviation increase in exposure is associated with a 0.08 p.p. increase in the CRE nonperforming rate (column (1)), which is more than 15 percent of the average CRE nonperforming rate of 0.5 percent. Our estimate is not

[10] We use net chargeoffs which include negative values from recoveries.

[11] Results are similar if we normalize provisions and chargeoffs by total assets instead of CRE loan volume.

driven by the possibility that more-exposed banks have inherently high CRE nonperforming rates. These regressions include the CRE nonperforming rate as of 2019Q4 as an additional control, and a placebo regression of the 2019Q4 CRE nonperforming rate on exposure results in a negative but statistically insignificant estimate for exposure. Columns (2) and (3) show that we find no significant association between exposure and chargeoffs or exposure and loan provisions although the point estimate for provisions is positive and close to being significant.[12]

Although we do find a sizable and significant effect of exposure on the nonperforming loan rate, exposure may create distress for banks in ways that are not fully captured by Table 4. The average CRE nonperforming rate was just 0.5 percent as of 2023Q4 even as national CRE prices had fallen 20 percent and vacancy rates had risen sharply by then. Delinquency rates and chargeoffs are often lagging indicators of distress. For example, a bank may recognize that a loan is unlikely to pay off in full due to a sharp decline in operating income at the property collateralizing the loan, but delinquency may not happen until the loan maturity date, as CRE loans are long, with low interim amortization payments and balloon payments due at maturity. Furthermore, operating income at the property may not have fallen yet but may be expected to fall in the future: office and retail leases are long (5-10 years), so landlords may still be collecting rents at pre-pandemic rates even if they expect rents on future leases to fall. Loss provisions are more forward-looking measures of distress. However, banks have some discretion in how much to allocate to loan loss provisions, and may have incentives to under-provision. Crosignani and Prazad (2024) provide evidence that large banks extended the maturity of impaired CRE loans to avoid writing-off those loans from the capital.[13] If such a practice is widespread even to smaller banks, our provisions measure will understate the true amount of distress in a bank's CRE portfolio.

[12] Because provisions and chargeoffs are measured as a total for all loan categories, we also tried including an interaction between exposure and bank assets, but we still did not find significant effects of exposure on chargeoffs or provisions.

[13] In addition, Hughes and Nichols (2025) provide evidence that banks can be slow to update loan-level internal risk assessments, which can affect loss provision requirements, in response to changes in credit risk.

Table 4: Effects of Exposure on Measures of Distress

	(1) Nonperf. Rate	(2) Chargeoffs	(3) Provisions
Exposure	0.0766**	-0.101	0.816
	(0.0337)	(0.569)	(0.539)
Bank Controls	Y	Y	Y
Dep. Variable Mean	.492	3.384	5.211
R-squared	0.0502	0.0459	0.392
N	1942	1942	1942

Standard errors in parentheses
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note: Nonperforming CRE loan rate is the dollar volume of CRE loans at least 90 days past due or in nonaccrual in 2023 Q4 divided by total CRE loans in 2023 Q4. Net chargeoffs is cumulative net chargeoffs (gross chargeoffs net of recoveries) across all types of loans between 2020 Q1 and 2023 Q4, normalized by the bank's CRE assets in 2019Q4. Provisions is loan loss provisions across all types of loans as of 2023 Q4, normalized by the bank's CRE assets in 2023Q4. All three dependent variables are multiplied by 100 so they are expressed in percentage terms. In addition to bank controls, the regressions shown in columns 1 and 3 include a control for the lagged dependent variable (measured in 2019Q4). Standard errors are heteroskedasticity robust.

4.2 Heterogeneous effects

4.2.1 Effects by county's WFHshare

Under our proposed mechanism, a bank's reduced CRE loan supply is driven by its incentives to de-risk its CRE loan portfolio. Thus, we would expect a larger decline of originations of riskier CRE loans. To test this hypothesis, we estimate a specification that includes an interaction of Exposure with the county's WFHshare, as defined in equation 2. Counties with high levels of WFHshare likely experienced worse declines in CRE fundamentals and are plausibly riskier counties for banks to lend in, especially because real estate prices tend to adjust slowly to changes in fundamentals. Column 4 in Table 5 presents estimates supporting the hypothesis. The point estimates imply that the marginal effect of exposure on credit supply is 5 percentage points more negative in a county with a one standard deviation higher level of WFHshare, all else equal. This result suggests that more exposed banks reduced credit supply by more in geographic areas that are plausibly more risky, consistent with bank de-risking in response to exposure of their loan portfolio to the WFH shock.

4.2.2 Effects by bank capitalization

Under our proposed mechanism, a bank with a larger CRE loan portfolio relative to its capital base should reduce credit supply by more in response to our exposure measure in the "Post" period (i.e. years 2020 and after), all else equal. Table 5 provides some support for this prediction. We interact Exposure with the size of the banks' CRE assets relative to Tier 1 capital, measured in 2019Q4 (pre-WFH shock). Column 2 shows that in the Post period, a bank with a high level of CRE assets relative to capital (95th percentile) lowers CRE lending by 1.1 percentage points more than a bank with the mean level of this ratio. However, the coefficient on the uninteracted measure of Exposure is still negative and significant, implying that even banks with very low levels of CRE assets relative to capital reduce CRE lending in the Post period. This result suggests that banks may target a specific level of CRE risk, in addition to an overall level risk in their loan portfolio.

Column 3 includes an additional set of interactions with the Tier 1 capital ratio, defined as Tier 1 capital relative to total bank assets, again measured in 2019Q4. The capital ratio measures the capital adequacy of the bank's entire balance sheet and is monitored by bank regulators. Consistent with exposure being a shock to bank expected net worth that causes them to de-risk, we find that the negative marginal effect of exposure for banks with high levels of CRE assets relative to capital is stronger (i.e. more negative) for banks with low capital ratios.

4.2.3 Effects by bank size

Column 5 of Table 5 shows the effects of Exposure on new CRE lending by bank size. The point estimates are fairly noisy since most banks in our sample are small and we cluster standard errors by bank. However, the point estimates show that the overall negative effect of Exposure on new CRE lending shown in column 1 is driven by banks in our sample with total assets less than $100 billion. For the 11 banks in our sample with over $100 billion in assets (the omitted category in the regression), the effect of Exposure on lending is actually slightly positive, but not statistically significant.

Why don't the very largest banks adjust their CRE credit supply in response to higher levels of exposure? One possibility is that the very largest banks are less risk averse because they can more easily cover realized losses from exposed CRE loans due to a low cost of raising external funds (Kashyap and Stein (1995)). In addition, large banks may be too big to fail and thus have a higher expected bailout probability, which may decrease their risk aversion (Naqvi and Pungaliya (2023)).

4.2.4 Spillovers: Effects of exposure on other lending

If more-exposed banks reduce CRE credit supply to derisk their balance sheet, then they may also reduce credit supply for other types of non-CRE, risky lending. In this section, we test for such spillover effects.

Our CoreLogic data do not cover non-real estate lending, so we use the Call Reports to test for spillover effects. Using these data, we measure the growth in a bank's outstanding loan amounts from 2019 to 2023 in C&I loans, consumer loans excluding residential mortgages, and residential mortgage loans. The Call Reports only contain data on outstanding loan balances, not new originations of credit, and do not include any geographic detail. We estimate bank-level regressions similar to equation 4, where Y_i refers to the centered growth rate of outstanding loan amounts from 2019 to 2023.

Table 6 presents the estimates. Before examining spillover to other loan types, we first

Table 5: Effects of Exposure on bank CRE purchase loan originations, various interactions

	(1) Loan Orig	(2) Loan Orig	(3) Loan Orig	(4) Loan Orig	(5) Loan Orig
Post x Exposure	-0.0728***	-0.0568***	-0.0699	0.302***	0.0693
	(0.0171)	(0.0214)	(0.0591)	(0.0680)	(0.0532)
Post x Exposure x CRE assets / Tier 1 Capital		-0.00678*	-0.0166**		-0.00601
		(0.00410)	(0.00770)		(0.00405)
Post x Exposure x CRE assets / Tier 1 Capital x Tier 1 Capital Ratio			0.0685*		
			(0.0350)		
Post x Exposure x WFHshare				-1.115***	
				(0.201)	
Post x Exposure x 10 - 100 billion in assets					-0.125
					(0.112)
Post x Exposure x 1 - 10 billion in assets					-0.138**
					(0.0599)
Post x Exposure x < 1 billion in assets					-0.0844
					(0.0537)
Bank Controls	Y	Y	Y	Y	Y
N	361032	361032	361032	361032	361032

Standard errors in parentheses
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: Poisson regression estimates. Post is a dummy variable equal to one if the origination year is 2020 or later. A complete set of exposure interaction terms for Post = 0 are included in each regression but are not shown. Standard errors are clustered by bank.

confirm that bank CRE outstanding loan growth based on Call Reports is also lower for more-exposed banks. Indeed, column 1 shows that for a bank with a mean level of CRE assets-to-Tier 1 Capital (1.64), a one standard deviation increase in exposure decreases the centered growth rate of CRE loans outstanding by 1 percentage points ($.0182 + 1.64 \times 0.585$). This effect is smaller in magnitude than the estimates presented in Section 3 using new originations of credit from CoreLogic, likely because reduced credit supply has a more indirect effect on outstanding loan balances compared to new originations of credit.

The second column provides evidence for spillover effects. For a bank with a mean level of CRE assets-to-Tier 1 Capital, a one standard deviation increase in exposure decreases the centered growth rate of non-CRE loans outstanding by 2 percentage points ($1.36 + 1.64 \times 0.421$), similar in magnitude to the results shown in column 1 for CRE lending.

Columns 3-5 show results separately for each loan category. There appear to be spillover effects for all three loan categories, though the estimates are not statistically significant for consumer and RRE lending. The spillover effects are strongest and significant for C&I lending. In addition, because the point estimates are consistently negative across the columns in Table 6, there is no evidence that more-exposed banks are derisking their CRE loan portfolio while increasing risk in other parts of their loan portfolio.

Table 6: Spillover effects

	(1) CRE	(2) Non-CRE	(3) CandI	(4) Consumer	(5) RRE
Exposure	-0.000182	-0.0136*	-0.0176	-0.0530***	-0.0359***
	(0.00922)	(0.00742)	(0.0119)	(0.0158)	(0.00915)
Exposure x CRE Assets-to-Tier 1 Capital	-0.00585**	-0.00421*	-0.00719**	-0.00339	-0.00153
	(0.00264)	(0.00255)	(0.00323)	(0.00435)	(0.00271)
Dep. variable mean	.383	.277	.228	.053	.269
Bank Controls	Y	Y	Y	Y	Y
R-squared	0.0668	0.0810	0.0352	0.0300	0.0920
N	1888	1890	1872	1868	1886

Standard errors in parentheses
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note: Growth is the centered growth rate of avg balance in 2023 relative to avg balance in 2019 measured using the call report balance variables. In addition to bank controls, results include a control for lagged growth in the dependent variable (2019 relative to 2015). Standard errors are heteroskedasticity robust.

5 Robustness and Additional Specifications

5.1 Separating Credit Supply from Credit Demand

For the main results in Section 3, we present the results of several alternative specifications that together alleviate concerns that our estimates pick up credit demand effects in addition to or instead of credit supply effects.

5.1.1 Leave-one-out measure of Exposure

We re-estimate the Poisson model in equation 6 using a leave-one-out measure of Exposure defined as:

$$Exposure_{ic} = \sum_{j \neq c} \frac{\text{CRE loans outstanding}_{ij}}{\text{total CRE loans outstanding}_{i}} \times WFHshare_{j} \tag{5}$$

where j indexes all counties in the country other than county c, i indexes the bank, and county-level $WFHshare$ is defined as in equation 2. Note that Exposure now varies by bank and county, whereas it varied only by bank in the baseline specification in Section 3. Leaving out county c from exposure means that exposure is defined without using any information from county c, making it even less likely that exposure is correlated with the bank's demand conditions in county c. The leave-one-out exposure is undefined for banks with CRE loans outstanding in only a single county, and these banks are dropped from our regressions that use leave-one-out exposure. Column 2 of Appendix Table 8 shows that our estimates are little changed when using the leave-one-out measure of exposure.

5.1.2 Within-county variation in CRE demand

Our main estimates include county-by-year fixed effects. Since our goal is to estimate credit supply effects and not credit demand effects, any within-county variation in CRE demand conditions that is correlated with Exposure could bias our results. For example, banks with higher exposure may tend to lend within areas of a county that are more likely to experience a negative CRE demand shock during the pandemic (e.g. more urban areas of the county). In addition, within a county, more exposed banks could be office lending specialists, and office demand declined disproportionately after 2020.

We address these concerns with three robustness checks. First, we rerun our main reduced-form specification at the zipcode level; i.e. the location index j in equation 6 represents a zipcode instead of a county and we include zipcode-by-year instead of county-by-year fixed effects.[14] Column 3 of Appendix Table 8 shows that our results are very similar to our main specification at the county-level.

Second, because there still may be heterogeneity in demand conditions within a zipcode, we include an additional bank-county level control variable for the downtown-ness of the bank's CRE portfolio within a county. This control variable addresses a potential concern that more exposed banks tend to lend in downtown areas, and downtown areas likely experienced a larger negative demand shock during the pandemic. We construct this control variable using our CoreLogic data. We first compute the distance of each CRE property in our data from the central business district (CBD) of its metropolitan area.[15] Then, for each bank-county pair (i, c), we compute the average distance-to-CBD of the CRE properties in i's loan portfolio in c. Column 4 of Appendix Table 8 presents results with this control variable added.[16] The results are very similar to the results from our main specification. The distance-to-CBD control variable has a positive effect on CRE lending, though it is not statistically significant.

Third, to address concerns about bank-property-type specialization, we estimate our main Poisson lending model at the bank-county-property type-year level. That is, we estimate

$$\lambda_{i,c,z,t} = \exp(\beta_t \text{Exposure}_i + X_i \gamma + \delta_{c,z,t}) \tag{6}$$

where z indexes the 8 different property types shown in Figure 6. Column 5 of Appendix Table 8 shows the results with the property-type fixed effects are very similar to the results from our main specification, suggesting that bank-property type specialization is not a main source of bias for our main results.

[14]The QCEW data are only available at the county level, not the zipcode level, and so our data do not allow for a zipcode-based measure of Exposure.

[15]We do this using the CoreLogic-provided latitude and longitude of the property. For micropolitan statistical areas, we do not have CBD coordinates and we set CBDness equal to missing for properties in such areas.

[16]We also include a dummry variable equal to one of the CBD variable is missing for the bank-county.

5.2 Measurement error in Exposure

One concern with our main estimates is that we measure banks' true exposure to the WFH shock with error. If so, this would tend to attenuate our estimates. In the section, we explore two potential sources of measurement error and conclude that they are unlikely to be large.

5.2.1 Alternative measure of the WFH shock

Our baseline results use a measure of exposure that relies on estimates of the work-from-home rate by industry as constructed by Dingel and Neiman (2020) (DN). As discussed in Section 3, the DN measure is based on pre-pandemic data and does not use any information on whether workers in a particular industry actually worked from home during the pandemic. Although these features of the DN measure help support the exogeneity of our exposure metric, they may also add to measurement error in our measure of exposure.

As an alternative to DN, we use estimates from the Current Population Survey (CPS) using post-pandemic data. Starting in October 2022, respondents to the CPS were asked if they teleworked or worked from home for pay at any time during the survey reference week. From October 2022 through December 2023, the share of respondents reporting "yes" to this question has been little changed at around 20 percent. The CPS also reports the industry in which the respondent is employed.[17] We compute the average work-from-home share for each industry k (WFH_k) using the share of respondents employed in industry k who report "yes" to the work-from-home question between October 2022 and December 2023. We then compute the county-level work-from-home share and $Exposure$ as in equations 2 and 1. Appendix Figure 8 shows a scatter plot of the DN and CPS-based measures of county-level work-from-home. The correlation coefficient between the two series is 0.7. Column 6 of Appendix Table 8 reproduces our main results using the CPS-based measure of work-from-home, which are very similar to our main results.

[17] The CPS uses census industry classification codes. We convert these to three-digit NAICs codes using the crosswalk provided by the CPS.

5.2.2 Exposure by property type

In our main specification, a bank's Exposure to the WFH shock is calculated by pooling all of the banks' CRE loans together, ignoring potential heterogeneity in the banks' exposure to different CRE property types. This may add measurement error to the exposure measure if a bank's true exposure depends on the property-type mix of its loan portfolio, causing us to understate the effects of exposure on credit supply.

To examine whether property-type heterogeneity is important, we consider property-type (p) specific exposure:

$$Exposure_{ip} = \sum_{c} \frac{\text{CRE loans outstanding}_{ipc}}{\text{total CRE loans outstanding}_{i}} \times WFHshare_c \tag{7}$$

where the sum of exposure across property types, $\sum_p Exposure_{ip}$, equals our baseline measure of Exposure defined in equation 1.

We estimate the following Poisson regression with the Exposure measures for different properties:

$$\lambda_{i,c,t} = \exp(\sum_{p} \beta_{0,p}\text{Exposure}_{i,p} + \beta_{1,p}\text{Exposure}_{i,p} \times Post_t) + \gamma X_i + \delta_{c,t}) \tag{8}$$

where $Post_t$ is an indicator variable for the post-pandemic 2020-2023.

Figure 6 presents the estimates of β_p for each property type. Exposure to office, multi-family, and "other" property types have relatively stronger effects on lending.[18] However, overall the point estimates across property types are fairly similar to each other and exposure to every property type is estimated to have a negative effect on lending (though the negative effect is not always statistically significant). Although there is scope for some mea-

[18] One might expect the office effect to be even more negative relative to the other property types given that the WFH shock most directly affected office properties. Our office category, however, includes many owner-occupied office properties, whose performance was much less adversely affected by WFH. (Unfortunately, we cannot identify owner-occupancy in our data.) In addition, the WFH shock to office demand had spillovers to other property types located in the same area. For example, retail stores near office buildings might have experienced a decline in their customers.

Figure 6: Effects of Exposure by property type on new bank CRE loan originations

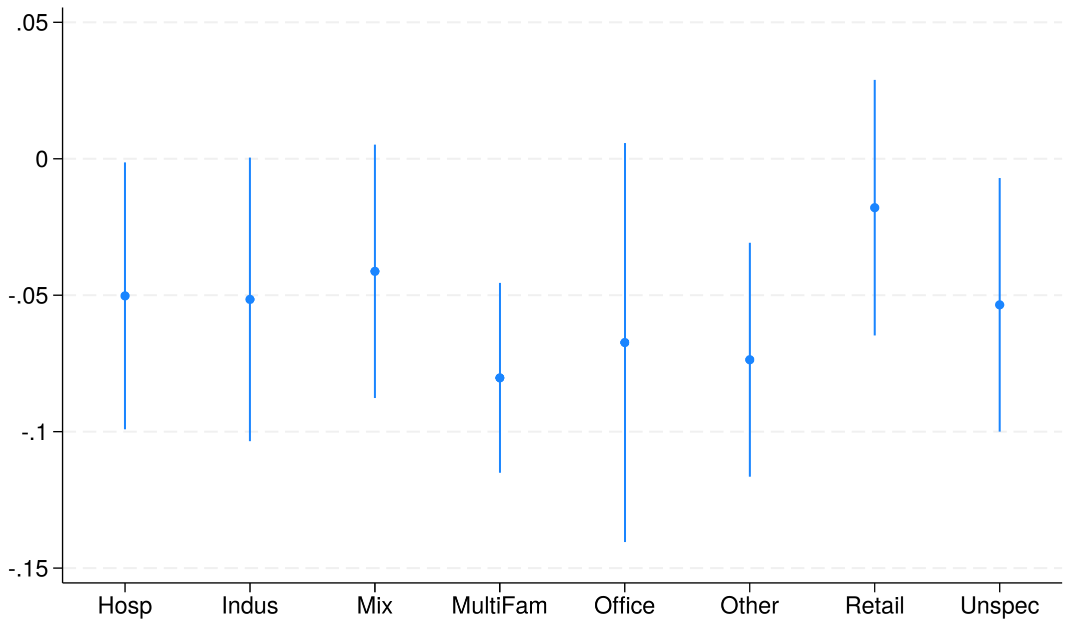

Note: This figure presents the estimates and the standard errors of $\beta_{1,p}$ in Equation 8 for each property type p. The standard errors are clustered by bank.

surement error in our exposure measure, the results in Figure 6 suggest that property-type heterogeneity is not a significant source of measurement error for our main results.

6 Effects on aggregate CRE lending

We have shown that more exposed banks reduced CRE lending by more. In this section, we examine whether the reduction in CRE credit supply by exposed banks was associated with a decrease in credit supply at the aggregate, market level. It is ex ante unclear whether banks' exposure to the WFH shock would lead to a decline in the aggregate credit supply because nonbanks are significant CRE lenders, and CRE loan borrowers might have substituted from banks to nonbanks when banks began to reduce their CRE lending.[19] However, there is important intermediary segmentation in CRE lending (Glancy et al. (2022a)) which may

[19] Nonbanks had a 40 percent market share of new CRE originations in the years right before and after the pandemic in our data, suggesting that they may be viable substitutes for banks.

act as a barrier to substitution across lender types. Indeed, we find that in markets where banks were on average more exposed to the WFH shock, the growth rate of overall CRE lending fell by more. The aggregate credit crunch is driven by banks: the average market experienced a 25 percentage point decrease in the overall growth rate of CRE originations, but a 56 percentage point decrease in the growth rate of new bank CRE originations as a result of bank exposure to the WFH shock. While nonbank lenders increase the pace of originations growth to fill this gap, they do not fully offset the reduction in CRE lending by banks.

Our CoreLogic data allows us to observe all CRE lending, from both banks and nonbanks such as insurance companies, mortgage REITs, pension funds, CMBS lenders, or others.[20] We exploit this feature of the data to test for aggregate effects using the cross-sectional regression

$$\frac{\text{Post-pandemic originations}_j - \text{Pre-pandemic originations}_j}{.5 \times (\text{Post-pandemic originations}_j + \text{Pre-pandemic originations}_j)} = \beta \text{BankExposure}_j + X_j \gamma + \epsilon_j \tag{9}$$

where j indexes the market, defined as a zipcode for these results. The dependent variable is zipcode-level centered growth rate in total CRE originations (bank + nonbank) in the post-pandemic period (2020-2023) relative to pre-pandemic period (2016-2019). We use the centered growth rate rather than a standard growth rate because it is bounded between -2 and 2, and so reduces the influence of outliers. We restrict the sample to zipcodes that have at least one origination in both the pre and post pandemic period. BankExposure for each zipcode is defined as

$$\text{BankExposure}_j = \frac{\sum_i \text{Exposure}_i \times \text{Pre-pandemic originations}_{i,j}}{\sum_i \text{Pre-pandemic originations}_{i,j}} \tag{10}$$

where Exposure is defined as in equation 1 and i indexes banks. BankExposure_j is the

[20]We do not, however, measure the exposure of these nonbanks to the WFH shock in the way we do for banks.

weighted-average exposure among banks that originate in zipcode j in the pre-pandemic period.[21] The distribution of BankExposure across zipcodes is shown in Figure 7. The mean and standard deviation are 0.33 and 0.03, respectively. If banks that are more exposed to the WFH shock reduce credit supply and that reduced supply cannot be offset by less exposed banks or nonbanks, then we would expect $\beta < 0$.[22]

For this specification, we define a market as a zipcode rather than a county as in Section 3.3. Observations in this specification are at the market level, not the bank-market level. We use a smaller definition of a market to increase sample size and provide precise estimates even once we include a large set of variables in X to control for differential demand conditions across markets (including state fixed effects).

Figure 7: Distribution of Bank Exposure Across Zipcodes

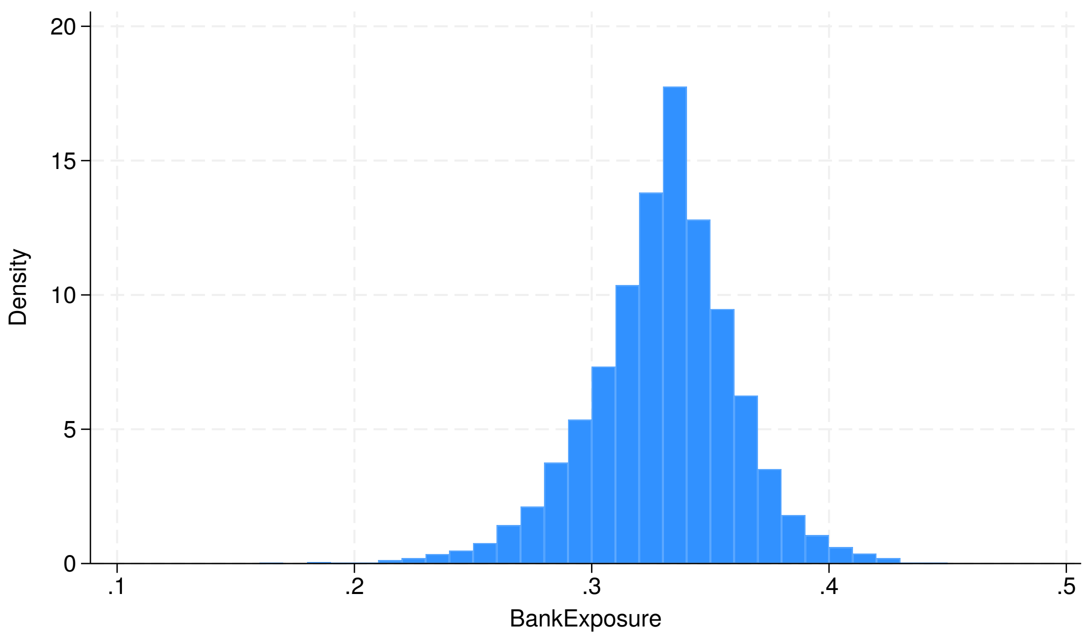

Note: Distribution of zipcode-level bank exposure, as defined in equation 10

X includes controls for other zipcode-level variables that may be correlated with the demand for CRE credit in the zipcode. For example, we include the actual work-from-home

[21]We drop zipcodes that have zero pre-pandemic bank originations.

[22]This assumes that bank-zipcode relationships are somewhat persistent so that the identities of banks lending in the pre-pandemic period is relevant for determining credit supply in the post-pandemic period.

share in 2023 and the share of population in the zipcode living in an urban area.[23] We also include the bank market share of new originations in the pre-pandemic period for each zipcode, calculated using the CoreLogic data. The full set of controls are listed in Table 7. Because the unit of observation is a market, we cannot flexibly control for changes in loan demand in a market by adding market fixed effects as we did in the bank-level regression in Equation 6. However, our control variables address some important threats to identification.

Table 7 reports the results. We estimate that bank exposure to the WFH shock induced by the pandemic is associated with a 25 percentage point decrease in the growth rate of total new CRE originations. We estimate this as -0.902*(0.33 - 0.05), where -0.902 is the coefficient on BankExposure in column 1, 0.33 is the average value of BankExposure across zipcodes, and 0.05 is the average pre-pandemic WFH rate as described in Section 3.3.

In columns 2 and 3, we decompose the effect of local bank exposure on aggregate credit supply into two parts: the effect of bank exposure on CRE originations by banks and the effect of bank exposure on CRE originations by nonbanks. Column 2 uses the growth rate in bank CRE originations in each zipcode as the dependent variable (instead of the growth rate of all zip code CRE originations). We estimate the average zipcode experienced a 56 percentage point decrease in the growth rate of new bank originations due to the BankExposure effect (-1.989*(0.33 - 0.05)). The average zipcode experienced 13 percentage points of bank lending growth (bottom panel of Table 7). Therefore, the BankExposure effect decreased the level of bank originations by about 43 percent (0.13 - 0.56). This estimate is comparable to the 47 percent decline implied by the bank-level estimates described in Section 3.3.

Our finding that bank exposure has a larger negative effect on bank origination growth (column 2) than on total origination growth (column 1) suggests the pull back in credit supply from banks was larger than for nonbanks in high bank-exposure counties. This result is shown more directly in column 3, where we redefine the dependent variable to be the

[23]These control variables come from the ACS. Using the 5-year 2023 survey, we calculate the actual WFH share using the share of workers aged 16 and over with at least $10,000 in earnings that reported "worked at home" as their usual mode of transportation to work.

growth rate in nonbank originations. Bank exposure is positively associated with nonbank lending growth, suggestive of some substitution from exposed banks to nonbanks. However, the estimate in column 1 for total loan growth is still negative and large, suggesting that this substitution was somewhat incomplete.

Coumns 4-6 show the estimates with a set of state fixed effects included. The standard errors get larger, but there is still evidence of large aggregate effects even using within state variation, and notwithstanding some substitution from exposed banks to nonbanks.

Table 7: Growth in zipcode-level CRE purchase originations

	(1) Total Orig.	(2) Bank Orig.	(3) NonBank Orig.	(4) Total Orig.	(5) Bank Orig.	(6) NonBank Orig.
BankExposure	-0.902***	-1.989***	0.509**	-0.722*	-1.432***	0.101
	(0.212)	(0.237)	(0.250)	(0.372)	(0.369)	(0.523)
Bank Orig. Share	-0.316***	-3.223***	0.727***	-0.310**	-3.575***	0.818***
	(0.100)	(0.110)	(0.149)	(0.152)	(0.175)	(0.218)
WFH share	0.0360	-0.691***	0.469***	0.225*	-0.207	0.354**
	(0.124)	(0.142)	(0.127)	(0.133)	(0.156)	(0.157)
Log(median income)	-0.205***	-0.0236	-0.226***	-0.174***	0.0229	-0.198***
	(0.0223)	(0.0246)	(0.0253)	(0.0344)	(0.0371)	(0.0416)
Urban share	-0.0459**	-0.0638**	-0.0829***	-0.0201	-0.0338	-0.0691
	(0.0227)	(0.0258)	(0.0275)	(0.0401)	(0.0529)	(0.0451)
College share	0.157**	0.382***	-0.100	0.0473	0.185**	-0.135
	(0.0646)	(0.0734)	(0.0722)	(0.0951)	(0.0824)	(0.110)
Log(population)	0.0727***	-0.00574	0.00796	0.0852***	0.0177	0.00613
	(0.00734)	(0.00849)	(0.00877)	(0.0150)	(0.0130)	(0.0198)
Change in Log(population)	0.0829*	-0.00265	0.0340	0.0712	-0.0129	0.00593
	(0.0451)	(0.0521)	(0.0577)	(0.0688)	(0.0785)	(0.0820)
Office share	0.125**	0.174***	0.0640	0.0779	0.166**	-0.00781
	(0.0503)	(0.0587)	(0.0587)	(0.0513)	(0.0694)	(0.0652)
Bank Orig. Share squared	0.783***	2.234***	0.764***	0.803***	2.439***	0.878***
	(0.0857)	(0.0948)	(0.157)	(0.121)	(0.146)	(0.224)
Dep. variable mean	.08	.13	.01	.08	.13	.01
stateFE	N	N	N	Y	Y	Y
R-squared	0.0727	0.115	0.169	0.107	0.170	0.206
N	11574	10405	9491	11574	10405	9491

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: Origination growth is the centered growth rate of total purchase originations in 2020-2023 relative to total purchase originations in 2016-2019 for each zipcode. All control variables are measured at the zipcode level and are based on pre-pandemic data, except for the WFH share, which is measured in 2023. Change in population is 2023 relative to 2019. Bank origination share is the share of total pre-pandemic originations that are bank originations, estimated using the CoreLogic data. Office share is the share of pre-pandemic originations that are associated with office properties, estimated using the CoreLogic data. Standard errors are heteroskedasticity robust.

7 Conclusion

We show that the WFH shock had large, negative effects on bank CRE credit supply. Our results are consistent with qualitative evidence of bank tightening over the pandemic period from the Senior Loan Officer Opinion Survey on Bank Lending Practices (SLOOS).[24] Our estimation approach allows us to quantify the magnitude of the tightening.

In addition, we find that the WFH shock had negative spillover effects on other risky bank lending and tighter bank CRE credit supply led to large, negative aggregate credit supply effects, accounting for possible substitution to nonbanks. These results suggest that the effects of the WFH shock on bank credit supply likely had important negative effects on the real economy. Looking at real outcomes such as new construction of CRE properties, directly, however, remains an important question for future research on the effects of the Covid-19 pandemic on the commercial real estate market.

References

Babalievsky, Fil, Kyle Herkenhoff, Lee E Ohanian, and Edward C. Prescott. 2023. "The Impact of Commercial Real Estate Regulations on U.S. Output."

Barrero, Jose Maria, Nicholas Bloom, and Steven J Davis. 2021. "Why working from home will stick."Technical report, National Bureau of Economic Research.

Barrero, José María, Nicholas Bloom, and Steven J. Davis. 2023. "The Evolution of Work from Home." *Journal of Economic Perspectives* 37 (4): 23–50. 10.1257/jep.37.4.23.

Bidder, Rhys M, John R Krainer, and Adam Hale Shapiro. 2021. "De-leveraging or de-risking? How banks cope with loss." *Review of economic dynamics* 39 100–127.

Buckman, Shelby R, Jose Maria Barrero, Nicholas Bloom, and Steven J Davis. 2025. "Measuring Work from Home."Technical report, National Bureau of Economic Research.

Büchler, Simon, Olivier Schöni, and Alex Van de Minne. 2024. "On the value of market signals: Evidence from commercial real estate redevelopment."

Chaney, Thomas, David Sraer, and David Thesmar. 2012. "The Collateral Channel: How Real Estate Shocks Affect Corporate Investment." *American Economic Review* 102 (6): 2381–2409. 10.1257/aer.102.6.2381.

[24] `https://www.federalreserve.gov/data/sloos.htm`

Chodorow-Reich, Gabriel. 2014. "The employment effects of credit market disruptions: Firm-level evidence from the 2008–9 financial crisis." *The Quarterly Journal of Economics* 129 (1): 1–59.

Crosignani, Matteo, and Saketh Prazad. 2024. "Extend-and-Pretend in the U.S. CRE Market."Technical report, Federal Reserve Bank of New York Staff Reports.

Davis, Morris A, Andra C Ghent, and Jesse Gregory. 2024. "The work-from-home technology boon and its consequences." *Review of Economic Studies* rdad114.

DeFusco, Anthony A, Charles G Nathanson, and Michael Reher. 2023. "Real Effects of Rollover Risk: Evidence from Hotels in Crisis."Technical report, National Bureau of Economic Research.

Delventhal, Matt, and Andrii Parkhomenko. 2023. "Spatial implications of telecommuting." *Available at SSRN 3746555.*

Delventhal, Matthew J, Eunjee Kwon, and Andrii Parkhomenko. 2022. "JUE Insight: How do cities change when we work from home?" *Journal of Urban Economics* 127 103331.

Dingel, Jonathan I, and Brent Neiman. 2020. "How Many Jobs Can be Done at Home?" *Journal of Public Economics* 189.

Froot, Kenneth A, and Jeremy C Stein. 1998. "Risk management, capital budgeting, and capital structure policy for financial institutions: an integrated approach." *Journal of financial economics* 47 (1): 55–82.

Ghent, Andra, and Rossen Valkanov. 2016. "Comparing Securitized and Balance Sheet Loans: Size Matters." *Management Science* 62 (10): 2784–2803. 10.1287/mnsc.2015.2260.

Glancy, David, John R. Krainer, Robert J. Kurtzman, and Joseph B. Nichols. 2022a. "Intermediary Segmentation in the Commercial Real Estate Market." *Journal of Money, Credit and Banking* 54 (7): 2029–2080. https://doi.org/10.1111/jmcb.12889.

Glancy, David, and Robert J. Kurtzman. 2024. "Determinants of Recent CRE Distress: Implications for the Banking Sector." *Finance and Economics Discussion Series.*

Glancy, David, Robert J. Kurtzman, and Lara Loewenstein. 2022b. "Loan Modifications and the Commercial Real Estate Market." *Finance and Economics Discussion Series* (2022-050): 1–71. 10.17016/FEDS.2022.050.

Glancy, David, Robert Kurtzman, Lara Loewenstein, and Joseph Nichols. 2023. "Recourse as shadow equity: Evidence from commercial real estate loans." *Real Estate Economics* 51 (5): 1108–1136. https://doi.org/10.1111/1540-6229.12450.

Glancy, David, and J. Christina Wang. 2024. "Lease Expirations and CRE Property Performance."Technical report, Federal Reserve Bank of Boston Research Department Working Papers.

Greenstone, Michael, Alexandre Mas, and Hoai-Luu Nguyen. 2020. “Do credit market shocks affect the real economy? Quasi-experimental evidence from the great recession and “normal” economic times.” *American Economic Journal: Economic Policy* 12 (1): 200–225.

Gupta, Arpit, Vrinda Mittal, and Stijn Van Nieuwerburgh. 2022. “Work From Home and the Office Real Estate Apocalypse.”Technical report, National Bureau of Economic Research.

Howard, Greg, Jack Liebersohn, and Adam Ozimek. 2023. “The short-and long-run effects of remote work on US housing markets.” *Journal of Financial Economics* 150 (1): 166–184.

Hughes, Samuel, and Joseph Nichols. 2025. “No News is Bad News: Monitoring, Risk, and Stale Financial Performance in Commercial Real Estate.”Technical report, Federal Reserve Board Working Paper.

Ivashina, Victoria, and David Scharfstein. 2010. “Bank lending during the financial crisis of 2008.” *Journal of Financial Economics* 97 (3): 319–338. https://doi.org/10.1016/j.jfineco.2009.12.001, The 2007-8 financial crisis: Lessons from corporate finance.

Jiang, Erica, Gregor Matvos, Tomasz Piskorski, and Amit Seru. 2023. “Monetary Tightening, Commercial Real Estate Distress, and US Bank Fragility.”

Kashyap, Anil K, and Jeremy C Stein. 1995. “The impact of monetary policy on bank balance sheets.” In *Carnegie-rochester conference series on public policy*, Volume 42. 151–195, Elsevier.

Khwaja, Asim Ijaz, and Atif Mian. 2008. “Tracing the impact of bank liquidity shocks: Evidence from an emerging market.” *American Economic Review* 98 (4): 1413–1442.

Mondragon, John A, and Johannes Wieland. 2022. “Housing demand and remote work.”Technical report, National Bureau of Economic Research.

Naqvi, Hassan, and Raunaq Pungaliya. 2023. “Bank size and the transmission of monetary policy: Revisiting the lending channel.” *Journal of Banking & Finance* 146 106688.

Peek, Joe, and Eric S Rosengren. 2000. “Collateral damage: Effects of the Japanese bank crisis on real activity in the United States.” *American Economic Review* 91 (1): 30–45.

Ramcharan, Rodney, Stephane Verani, and Skander J Van den Heuvel. 2016. “From Wall Street to main street: the impact of the financial crisis on consumer credit supply.” *The Journal of finance* 71 (3): 1323–1356.

Rosenthal, Stuart S., William C. Strange, and Joaquin 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. https://doi.org/10.1016/j.jue.2021.103381, JUE Insights: COVID-19 and Cities.

A Details on the match between CoreLogic and bank regulatory filings

We match CoreLogic lender names to RSSD IDs using the following algorithm:

1. **Name cleaning**: we clean the names of CoreLogic lenders and the names of institutions with RSSD IDs. This involves removing punctuation, removing auxiliary words like "national association", "NA", "INC", etc. We also replace common words with their abbreviations, for example, replacing "bank" with "bk" and replacing "California" with "CA".

   One of the most common punctuation marks used in CoreLogic lender names is "/ ", which seems to have two different meanings depending on context. In some cases, "/" seems to refer branch locations of a bank. For example, names such as "Wells Fargo BK/AZ" and "Wells Fargo BK/TX" seem to point to particular Wells Fargo branch locations in Arizona and Texas. In other cases, "/" seems to be used in place of "of" in names such as "First State BK/ Gainesville". Because it is difficult to know what "/" means without investigating each specific case, we initially replace "/" with "of" in all cases at the cleaning stage. Although this change is not correct for all names containing "/", we make sure to assign correct RSSD IDs for examples like "Wells Fargo/AZ".

2. **Exact name match**: we assign RSSD IDs to CoreLogic lenders by finding CoreLogic-RSSD pairs that have the exactly clean names in both data sets.

3. **One name includes the other as a substring**: among CoreLogic lenders that fail to match exactly to any RSSD names, we find potential RSSD IDs for these CoreLogic lenders by finding CoreLogic-RSSD pairs such that the shorter name of a pair is included in the longer name of the pair. For example, the RSSD ID of "Wells Fargo BK" will be matched to CoreLogic lender "Wells Fargo BK/TX" in this step.

4. **Identifying duplicate matches**: there are many legally distinct banks with different RSSD IDs, but with the exactly same names like "Centennial BK" or nearly identical names like "Centennial Bank and Trust". Thus, we identify cases where a single CoreLogic lender name is matched to multiple RSSD IDs and address this issue at later steps.

5. **Use lender location information**: CoreLogic provides information on lender addresses, and this field is populated for roughly half of observations. Whenever a CoreLogic lender name is matched to multiple RSSD IDs and has its location information, we choose the RSSD ID with at least one branch at the same zipcode (or the same city if the zipcode information is missing). We use the Summary of Deposits (SOD) for the information on bank branch locations. Among duplicate matches, we classify a CoreLogic-RSSD pair as matched only if there is only one RSSD ID with branches in a particular zipcode.

6. **Use property location information**: Using matched pairs from the previous steps, we can identify which RSSD IDs typically lend against properties in certain locations (zipcode, city, or state). This allows us to drop false matches when multiple RSSD IDs are matched to the same Corelogic name, and ensure that the loan is assigned to the most likely RSSD ID. For example, we identify that Centennial BK's lending against properties in Oregon is exclusively done by only RSSD ID XXX although there are multiple RSSD IDs with names identical or almost identical to Centennial BK. Then, whenever we see Centennial BK's lending against other properties in Oregon with the lender address missing, we assign XXX as that loan's RSSD ID.

After these steps, we match about 91% of bank CRE loans in CoreLogic to unique RSSD IDs in term of the number of loans. 70% of bank CRE loans are matched with the exact name match (step 2); 11% are matched using lender addresses (step 5); 8.7% are matched using property addresses (step 6); and less than 1% of loans are uniquely matched using the shorter name match (step 3).

In terms of loan amount, the overall match rate is about 99 percent among bank CRE loans, and 94 percent of bank CRE loan amounts are matched with the exact name match (step 2).

B Additional empirical results

Our main specification in Section 3 uses the count of purchase originations as the dependent variable. Column 8 of Appendix Table 8 shows results when the dependent variable is the dollar volume of purchase originations. Overall, the results are similar to the baseline results (column1). The point estimates are a little more negative using dollar volumes, though they are more imprecisely estimated too.

Column 7 of Appendix Table 8 shows results for the count of new refinance originations. Interestingly, there is no evidence for an effect of bank exposure on refinance credit supply. This could be because if banks reduce refinance credit supply to their existing borrowers, then they may end up creating more losses for themselves, as existing borrowers who are unable to refinance balloon payments due at loan maturity may simply default.

Figure 8: WFH share by county based on Dingel-Neiman and CPS

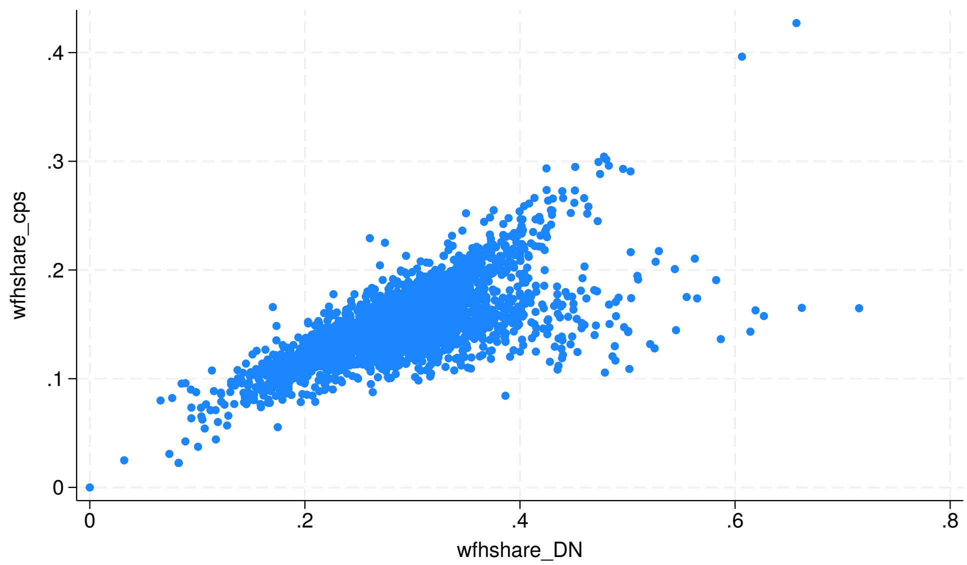

Table 8: Robustness and Alternative Specifications

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Baseline	lvo	Zip. FE	CBDness	Prop. FE	CPS	Refi	Dollars
Post x Exposure	-0.0728***	-0.0738***	-0.0736***	-0.0732***	-0.0671***	-0.0650***	-0.0158	-0.107***
	(0.0171)	(0.0136)	(0.0133)	(0.0171)	(0.0175)	(0.0179)	(0.0148)	(0.0318)
CBDness				0.0269				
				(0.0202)				
Bank Controls	Y	Y	Y	Y	Y	Y	Y	Y
N	361032	271072	2422972	361032	707396	361032	378552	360669

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: Baseline reproduces the main estimates from column 1 of Table 5. The remaining columns are variants of the baseline specification, where "lvo" denotes the leave-one-out measure of exposure; "Zip. FE" denotes zipcode-by-year fixed effects instead of county-by-year fixed effects as in column 1; "CBDness" denotes a bank-county control variable added for the CBDness (measured in hundreds of miles) of the bank's loan portfolio within a county; "Prop. FE" denotes property type-by-county-by-year fixed effects; "CPS" denotes the Current Population Survey-based measure of Exposure; "Refi" denotes Refinance loan originations as the dependent variable; "Dollars" denotes that the dependent variable is dollars of new originations, rather than counts.