When Oil Narratives Become Bank Risk

Gabriel A. Abeyie*
University of North Carolina Wilmington

June 3, 2026

Abstract

Oil-market news can signal very different risks for banks exposed to oil-producing regions. This paper studies how newspaper-based oil narratives are associated with U.S. bank balance sheets. Supply-disruption narratives predict higher cumulative loan-loss provisions and net charge-offs, lower loan growth, lower deposit growth, and weaker short-run profitability among more oil-exposed banks. Demand-salience narratives move in the opposite direction, predicting stronger lending, deposit growth, and profitability, while provisions and charge-offs decline. These findings show that oil narratives do not have one banking effect: some oil stories become bank risk, while others coincide with expansion in oil-exposed banking markets.

Keywords: oil narratives; bank risk; local banks; oil shocks; text analysis; financial stability

JEL Classification: G21, G28, E44, Q41, C23

*Cameron School of Business, Department of Economics and Finance, University of North Carolina Wilmington. Email: abeyieg@uncw.edu

1 Introduction

Oil shocks do not have one banking effect. A disruption to supply, a collapse in demand, an inventory glut, and a geopolitical scare can all raise oil-market attention, but they do not imply the same local balance-sheet response. For banks exposed to oil-intensive counties, some oil narratives signal borrower stress and future credit losses. Others coincide with higher regional income, stronger loan demand, and deposit growth. Treating oil news as a single risk signal therefore misses the main economic distinction: the banking effect depends on which oil story becomes salient.

This paper studies when oil narratives become bank risk. I combine newspaper-based measures of oil narratives with quarterly U.S. bank balance-sheet data from 2001 to 2024. I measure bank exposure to oil-producing regions using branch-level deposit locations from the FDIC Summary of Deposits and county oil-and-gas payroll intensity from County Business Patterns. The empirical design estimates dynamic panel local projections (Jordà, 2005) that compare banks with different degrees of oil exposure within the same quarter. Bank fixed effects absorb persistent differences across institutions. Quarter fixed effects absorb national banking conditions, monetary policy, aggregate oil prices, and other common shocks. The identifying variation comes from whether more oil-exposed banks respond differently when a particular oil narrative becomes more prominent.

The central result is that oil narratives have sharply different banking effects. Supply-disruption narratives behave like adverse credit-risk signals for banks with greater exposure to oil-intensive counties. They predict higher cumulative loan-loss provisions and higher cumulative net charge-offs. They also predict lower cumulative loan growth, lower deposit growth, and weaker profitability. By contrast, demand-salience narratives move in the opposite direction: they predict stronger lending, stronger deposit growth, and higher profitability, while provisions and charge-offs fall. Aggregate oil-risk attention hides these offsetting channels.

The distinction is large enough to matter economically. Over an eight-quarter horizon, supply-disruption narratives raise cumulative provisions/assets and charge-offs/assets by 0.11 and 0.12 percentage points respectively for a 10 percentage-point increase in deposit-weighted oil exposure. They also reduce cumulative loan growth by 1.81 percentage points and deposit growth by 1.97 percentage points. The corresponding demand-salience coefficients have the opposite signs: loan growth rises by 2.39 percentage points, deposit growth rises by 1.72 percentage points, and provisions and charge-offs decline. These effects are estimated after allowing more oil-exposed banks to respond differently to oil returns, realized oil-price volatility, and oil-price-pressure tone.

The results are not a Texas-only pattern. Texas is central to U.S. oil production[1] and banking exposure, so a natural concern is that the findings are driven entirely by Texas counties. I address this by reconstructing the exposure measure after setting Texas counties' oil intensity to zero while keeping Texas branch deposits in the denominator. The credit-loss channel survives: supply-disruption narratives continue to predict higher provisions and charge-offs among banks exposed to non-Texas oil-intensive counties. The lending contraction weakens, indicating that Texas contributes importantly to the loan-growth response, but the main credit-risk result is not driven only by Texas county exposure.

The paper builds on the insight that not all oil shocks are alike. The oil-macro literature has long emphasized that supply shocks, demand shocks, and precautionary-demand shocks have different aggregate effects (Kilian, 2009; Baumeister and Hamilton, 2019). This paper applies the same logic to narratives and banks. The relevant object is not only the oil price or the amount of oil news. It is the type of oil narrative. A supply-disruption story is not the same banking shock as a demand-salience story, even if both raise oil-market attention.

The mechanism is straightforward. Supply-disruption narratives capture news about outages, shortages, refinery constraints, production interruptions, hurricanes, sanctions, and other disruptions to oil-market activity. For banks with deposit bases located in oil-intensive counties, these narratives can signal stress among energy bor-

[1] The U.S. Energy Information Administration (EIA) identifies Texas as the top crude-oil- and natural-gas-producing state and the state with the most refining capacity in the United States.

rowers, firms connected to local oil production, and households in oil-dependent regions. The balance-sheet response is therefore visible in credit-loss recognition and slower growth. Demand-salience narratives, by contrast, capture periods when oil demand, global growth, consumption, and macroeconomic activity become prominent in oil-market coverage. In the data, these narratives are associated with stronger lending, deposits, and profitability among more oil-exposed banks, suggesting that they often coincide with expansionary local conditions rather than borrower distress.

The contribution is threefold. First, the paper connects oil narratives to bank balance-sheet risk. Recent work uses newspaper text to study bank lending and loss recognition, including general media sentiment (Agoraki et al., 2022) and energy-related uncertainty (Pacelli et al., 2025), while a separate literature studies how oil-price shocks affect banks exposed to oil-producing regions (Wang, 2021; Gilje, 2019). To my knowledge, this paper is the first to separate oil narratives by type and ask whether supply-disruption and demand-salience narratives predict different balance-sheet outcomes across banks with different local exposure to oil-intensive counties. The relevant object is not the volume of oil news but the type of oil story facing exposed banks. Second, the paper shows that aggregate oil-risk attention is too coarse for banking applications. The same aggregate index can combine narratives that have opposite balance-sheet implications. Third, the paper provides evidence that oil narratives are associated with local banking exposure measured from branch-deposit locations and county oil-and-gas intensity.

The results also speak to financial-stability measurement. Banks do not only respond to realized prices and realized defaults. They operate in information environments where news about disruptions, demand conditions, and market stress shapes expectations about borrowers and local activity. Newspaper narratives provide a real-time measure of that information environment. The evidence here suggests that narrative type matters: some oil stories become credit risk, while others accompany balance-sheet expansion.

The rest of the paper proceeds as follows. Section 2 places the paper in the literature. Section 3 describes the oil narrative measures, bank balance-sheet data, and oil exposure measure. Section 4 presents the empirical specification. Section 5 reports the aggregate oil-risk results and the narrative-type decomposition. Section 6 presents robustness checks, including oil-price controls, the exclusion of Texas counties from the exposure construction, and an alternative headquarters-state exposure definition. Section 7 concludes.

2 Related Literature

This paper links three literatures: oil-shock heterogeneity, narrative measurement from text, and the transmission of local real shocks to banks.

The first is the macroeconomics of oil shocks. Hamilton (1983) established the central role of oil-price movements in U.S. business-cycle fluctuations, showing that oil-price increases preceded many postwar recessions. Subsequent work showed that the economic effects of oil-price movements depend on the source of the shock. Kilian (2009) decomposes oil-price movements into oil supply shocks, aggregate demand shocks, and oil-specific demand shocks, and shows that these shocks have different implications for real activity and prices. Baumeister and Hamilton (2019) further refine the identification of oil supply and demand components, while Känzig (2021) identifies oil supply news shocks using OPEC announcement surprises. The central lesson from this literature is that oil-price movements cannot be treated as a single economic object.

This paper brings that logic to oil narratives and banks. I do not identify structural oil supply and demand shocks in the sense of Kilian (2009). Instead, I measure the salience of different oil stories in newspaper coverage and ask whether those stories have different balance-sheet implications for banks exposed to oil-intensive counties. The distinction matters. A supply-disruption narrative can signal outages, production interruptions, sanctions, refinery constraints, or other disruptions to oil-market activity. A demand-salience narrative can instead reflect stronger oil demand, global growth, China demand, consumption, or broader macroeconomic conditions. The paper shows that these narrative types are associated with different balance-sheet responses among exposed banks.

The second related literature measures economic narratives and information from text. Shiller (2017) emphasizes that narratives spread through the economy and can shape economic fluctuations. Gentzkow et al. (2019) survey methods for converting text into economic data. Newspaper-based measures have become standard in macro-finance, with the economic policy uncertainty index of Baker et al. (2016) providing a leading example. Recent work also uses oil-news narratives to measure news-driven sentiment, uncertainty, and macroeconomic oil shocks (Abeyie, 2025). This literature uses media coverage and narrative classifications to measure attention, sentiment, uncertainty, and the perceived sources of economic fluctuations rather than relying on prices alone.

Within this literature, a large body of work measures financial-market information from news text. Tetlock (2007) shows that media tone predicts stock-market returns and trading activity. García (2013) finds that news sentiment is especially informative during recessions. Loughran and McDonald (2011) show that financial text requires domain-specific dictionaries, which informs how I count narrative and tone language in oil-news coverage. Manela and Moreira (2017) construct a newspaper-based measure of implied volatility from front-page coverage and show that news text recovers a risk signal distinct from prices. These studies establish that newspaper text carries financial information beyond contemporaneous price movements, which is the premise behind my use of oil narratives as a bank-relevant signal.

A smaller literature applies text-based measures directly to banks. Agoraki et al. (2022) use news-based sentiment to study U.S. bank lending and financial stability and find that lending falls when sentiment is weak, with the effect stronger for banks holding higher credit risk. Bae et al. (2024) extract sentiment from bank disclosures and show that more negative sentiment predicts higher loan-loss provisions and lower future lending. Relatedly, Ng et al. (2020) find that banks raise loan-loss provisions when policy uncertainty is high, consistent with provisions signaling expected losses. In the energy domain, Pacelli et al. (2025) find that energy-uncertainty shocks raise bank credit risk in the Eurozone, and Dell'Atti et al. (2025) link oil-and-gas price bubbles to banking risk. These studies rely on aggregate sentiment, aggregate policy or energy uncertainty, or managerial sentiment drawn from bank disclosures. They do not separate oil narratives by type or interact narrative salience with banks' local exposure to oil-intensive counties.

I contribute to this literature by moving from aggregate oil-news measurement and showing that the type of narrative matters for financial transmission. The paper does not use newspaper coverage only to construct an aggregate attention or uncertainty index in the spirit of Baker et al. (2016). Instead, it separates oil narratives into economically interpretable categories and shows that these categories have opposite implications for banks with greater exposure to oil-intensive counties. The novelty is not simply the use of text data; it is the combination of narrative-type measurement with bank balance sheets and local oil exposure. Aggregate oil-risk attention is informative but too coarse: it pools supply-disruption narratives that predict credit-loss recognition with demand-salience narratives that coincide with stronger lending, deposits, and profitability.

The third literature studies how local real shocks transmit through banks. A central finding is that banks are not purely passive conduits of local conditions: branch networks integrate otherwise segmented regional markets, so that shocks in one area reach lending and deposits elsewhere. Becker (2007) documents geographic segmentation of U.S. capital markets, establishing that local deposit supply shapes local lending. Morgan et al. (2004) show that bank integration alters how state business cycles comove, and Landier et al. (2017) show that the geographic integration of the banking market transmits local shocks into house-price comovement across states. Gilje et al. (2016) use shale-driven deposit windfalls to show that branch networks integrate U.S. lending markets, and Gilje (2019) shows that these local credit-supply shocks have real effects that depend on local banking-market structure. Cortés and Strahan (2017) show that financially integrated banks reallocate capital in response to local natural disasters. Closest to this paper, Wang (2021) studies banks operating in oil-concentrated counties and shows that oil-price declines generate deposit losses, credit-line drawdowns, troubled-loan increases, and lending contractions, with effects propagating through bank networks. The premise that branch-deposit geography is informative about a bank's local real exposure—which underlies my exposure measure—follows directly from this body of work.

My analysis differs from this work in its object of interest. Existing studies emphasize realized local shocks, oil-price movements, and bank exposure to oil-producing areas. I ask whether newspaper narratives about different types of oil-market episodes predict differential responses among exposed banks, even after controlling for oil returns, oil-price volatility, and oil-price-pressure tone. The empirical design uses bank and quarter fixed effects, so the estimates compare banks with different degrees of oil exposure within the same quarter when a particular oil narrative becomes salient.

The paper therefore joins the oil-shock-heterogeneity literature with the banking-transmission literature through narrative measurement. The main result is a narrative analogue of the "not all oil shocks are alike" principle: not all oil narratives become bank risk. Supply-disruption narratives predict higher provisions, higher charge-offs, and weaker balance-sheet growth among more oil-exposed banks. Demand-salience narratives instead predict stronger lending, deposit growth, and profitability. Distinguishing the type of oil story is therefore essential for understanding when oil-market information becomes bank risk.

3 Oil Narratives and Bank Exposure

This section describes the three empirical inputs used in the analysis: newspaper-based oil narrative measures, bank balance-sheet outcomes, and deposit-weighted oil exposure. The sample used in the bank regressions runs from 2001Q1 through 2024Q4, reflecting the availability of the quarterly bank panel.

3.1 Newspaper-Based Oil Narrative Measures

I construct monthly oil narrative measures from a corpus of oil-related articles in major U.S. newspapers. The corpus is designed to capture national oil-market coverage rather than firm-specific press releases or local energy stories. Articles are selected using oil-market terms such as crude oil, oil prices, OPEC, Brent, WTI, petroleum, refineries, oil supply, oil demand, inventories, sanctions, embargoes, hurricanes, conflicts, and related market-relevant language. Appendix A provides additional details on the construction of the text measures, and Appendix Table 5 summarizes the main indices and oil-market controls used in the empirical analysis.

The main analysis focuses on two narrative types: supply-disruption narratives and demand-salience narratives. This choice is deliberate. The aggregate oil-risk index combines multiple oil stories that may have different implications for banks. A disruption narrative can signal stress for borrowers and firms tied to oil production, transportation, refining, or related local activity. A demand-salience narrative can instead reflect stronger oil demand, global growth, consumption, demand recovery, or broader macroeconomic conditions. Separating the two allows the empirical analysis to ask whether different oil stories have different balance-sheet implications for banks exposed to oil-intensive counties.

For each article, I count narrative-specific terms in oil-centered text windows. Supply-disruption narratives are based on language related to disruptions, shortages, outages, shutdowns, refinery problems, storms, sanctions, embargoes, supply cuts, and production interruptions. Demand-salience narratives are based on language related to demand, consumption, growth, China demand, global demand, economic activity, and demand slowdowns or recoveries. I aggregate article-level counts to the monthly frequency and transform each series using $\log(1 + x)$. The resulting monthly indices are standardized over the full text sample.

Figure 1 plots the two narrative indices. The series move together during some major oil episodes, but they are not the same object. Supply-disruption narratives spike when oil coverage centers on interruptions to supply, refining, sanctions, storms, or production constraints. Demand-salience narratives become prominent when oil coverage emphasizes oil demand, growth, consumption, recovery, and macroeconomic conditions. I therefore interpret demand salience as a measure of demand-related oil coverage, not as a structural adverse demand shock.

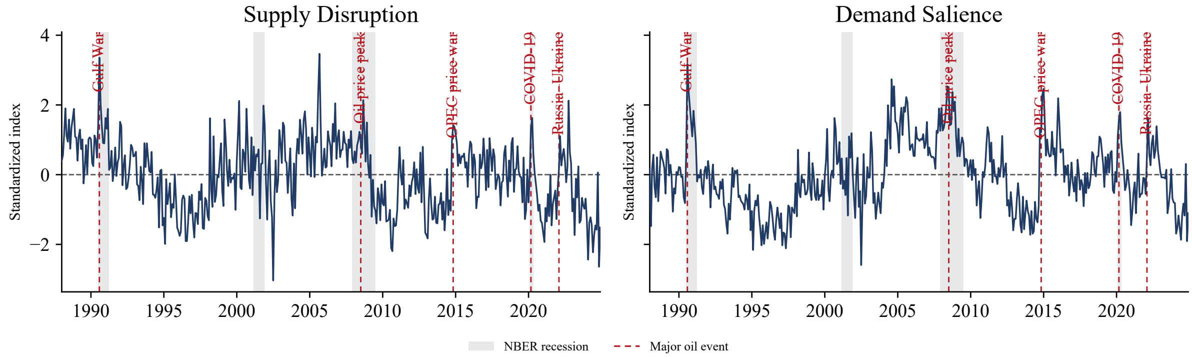

Figure 1: Demand and Supply Oil Narrative Indices
Notes: The figure plots standardized monthly indices for supply-disruption narratives and demand-salience narratives constructed from the oil-news corpus. Gray shaded areas denote NBER recessions.

The monthly narrative measures are converted to quarterly frequency by averaging the monthly standardized indices within each quarter. This timing matches the quarterly Call Report data. In the timing-clean baseline specification, the quarter-t narrative index is interacted with lagged deposit-weighted oil exposure to predict future bank outcomes. Flow outcomes are accumulated from quarter $t+1$ through horizon h, excluding quarter t, and stock outcomes are measured as changes from quarter $t-1$ to quarter $t+h$.

3.2 Bank Balance-Sheet Data

Bank balance-sheet data come from quarterly U.S. bank regulatory filings. The panel contains bank-level outcomes and controls from 2001Q1 through 2024Q4. The main outcomes are nonperforming loans, loan-loss provisions, net charge-offs, loan growth, deposit growth, return on assets, and equity/assets. These variables capture credit quality, loss recognition, balance-sheet growth, profitability, and capital.

The main credit-risk outcomes are measured as follows. The nonperforming-loan ratio is nonperforming loans divided by total loans. Loan-loss provisions and net charge-offs are scaled by total assets. Loan growth and deposit growth are measured at the quarterly bank level. Return on assets is quarterly net income divided by total assets. Equity/assets is total equity divided by total assets.

I winsorize the main bank ratios by quarter at the 1st and 99th percentiles to limit the influence of extreme reporting values. The baseline controls include log assets, equity/assets, loans/assets, deposits/assets, and return on assets. These controls account for size, capitalization, balance-sheet composition, funding structure, and recent profitability.

The matched bank-narrative panel contains 383,365 bank-quarter observations from 2001Q1 through 2024Q4. The panel is unbalanced because banks enter, exit, merge, or stop reporting over time. After merging lagged deposit-weighted oil exposure, roughly 95 percent of bank-quarter observations have nonmissing exposure information. Exact sample sizes vary slightly across outcomes and horizons because the local projections use future leads and cumulative future outcomes.

3.3 Oil Exposure

The baseline exposure measure uses banks' branch-deposit locations to measure exposure to oil-intensive counties. This measure is constructed from two public data sources. First, I use FDIC Summary of Deposits data to observe each bank's branch deposits by county. Second, I use County Business Patterns to measure county oil-and-gas intensity using payroll in oil-and-gas-related industries.

For county c in year y, I define oil intensity as the oil-and-gas payroll share:

$$OilIntensity_{c,y} = \frac{OilGasPayroll_{c,y}}{TotalPayroll_{c,y}}. \tag{1}$$

Oil-and-gas payroll includes oil and gas extraction, drilling oil and gas wells, and support activities for oil and gas operations. This measure captures the local importance of oil-related activity in a county's private-sector payroll.

For bank i, I compute the share of its deposits located in county c in year y:

$$DepositShare_{i,c,y} = \frac{Deposits_{i,c,y}}{\sum_{c'} Deposits_{i,c',y}}. \tag{2}$$

The baseline bank-level oil exposure measure is the deposit-weighted average of county oil intensity:

$$OilExposure_{i,y} = \sum_{c} DepositShare_{i,c,y} \times OilIntensity_{c,y}. \tag{3}$$

This measure captures the extent to which a bank's deposit base is located in counties where oil-and-gas activity is economically important; using branch-deposit geography to measure a bank's local real exposure follows Becker (2007) and Gilje et al. (2016). It is more granular than a headquarters-state indicator because it uses the geographic distribution of branch deposits rather than assigning the entire bank to a single state. A bank headquartered outside an oil-producing state can still be exposed if it operates branches in oil-intensive counties, and a bank headquartered in an oil-producing state can have low exposure if its deposits are located outside oil-intensive areas.

In the regressions, I use lagged exposure, $OilExposure_{i,t-1}$, measured using the prior calendar year's Summary of Deposits and County Business Patterns data. This timing avoids using branch or county information realized after the outcome quarter. The coefficients are scaled so that the reported estimates correspond to a 10 percentage-point increase in oil exposure.

Figure 2 summarizes the resulting exposure measure. The left panel shows that deposit-weighted oil exposure is highly skewed: many banks have little or no exposure to oil-intensive counties, while a smaller group has substantial exposure. This skewness is useful for identification because the estimates compare banks with different degrees of local oil exposure within the same quarter. The right panel shows that the upper tail of exposure remains economically meaningful throughout the sample, indicating that the results are not driven by a short-lived set of oil-exposed banks.

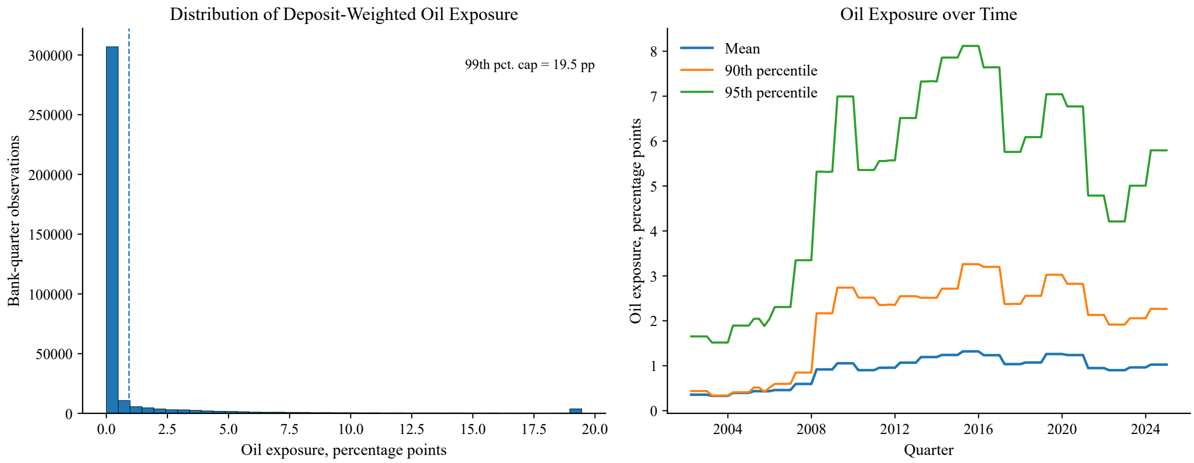

Figure 2: Deposit-Weighted Oil Exposure
Notes: The figure summarizes lagged deposit-weighted oil exposure. Oil exposure is constructed from FDIC Summary of Deposits branch locations and county oil-and-gas payroll intensity from County Business Patterns. The left panel plots the distribution of bank-quarter exposure, capped at the 99th percentile for readability. The right panel plots the mean, 90th percentile, and 95th percentile of exposure over time.

Table 1 provides benchmarks for the main variables used in the analysis. The bank outcomes are economically small in levels but vary meaningfully across banks and quarters: average quarterly provisions/assets and charge-offs/assets are about 0.05 percentage points, while average quarterly loan and deposit growth are about 1.6 percent. The exposure measure is highly skewed. The median bank-quarter has no exposure to oil-intensive counties, while the mean exposure is positive because a smaller group of banks has substantial branch-deposit exposure to oil-intensive local economies. This skewness is useful for identification because the local projections compare banks with different degrees of oil exposure within the same quarter. The narrative indices are standardized, so the reported coefficients can be interpreted as responses to a one-standard-deviation increase in narrative salience interacted with a 10 percentage-point increase in oil exposure.

Because deposit-weighted oil exposure is highly skewed, Appendix Table 8 reports the empirical support for positive and high-exposure observations. About one-third of bank-quarter observations have positive exposure, and nearly half of banks are exposed to oil-intensive counties at some point in the sample. The 10 percentage-point scaling used in the baseline tables is an upper-tail comparison, but it is not driven by a handful of banks: 9,427 bank-quarter observations and 238 banks reach exposure of

Table 1: Summary Statistics

Variable	N	Mean	SD	P10	P25	Median	P75
Bank outcomes							
NPL ratio	357,672	0.30	0.79	0.00	0.00	0.00	0.19
Provisions/assets	359,440	0.05	0.13	0.00	0.00	0.02	0.05
Charge-offs/assets	359,440	0.05	0.13	0.00	0.00	0.01	0.04
Loan growth	359,607	1.62	5.13	-3.26	-0.89	1.21	3.58
Deposit growth	359,607	1.59	5.45	-3.69	-1.29	0.99	3.69
ROA	359,440	0.25	0.27	0.05	0.15	0.25	0.35
Equity/assets	359,443	11.51	6.18	7.80	8.89	10.36	12.45
Exposure and narratives							
Deposit-weighted oil exposure	359,607	0.94	3.70	0.00	0.00	0.00	0.01
Oil exposure: top-5 county share	359,607	5.32	20.42	0.00	0.00	0.00	0.00
Oil exposure: top-10 county share	359,607	32.72	45.68	0.00	0.00	0.00	100.00
Aggregate oil-risk attention	359,607	-0.08	0.93	-1.22	-0.85	-0.08	0.62
Supply-disruption narratives	359,607	-0.10	0.84	-1.32	-0.79	-0.05	0.57
Demand-salience narratives	359,607	0.19	0.90	-1.01	-0.44	0.19	0.82
Oil-market controls							
WTI return	359,607	1.96	26.20	-19.36	-6.16	3.22	13.26
WTI realized volatility	359,607	4.47	0.71	3.63	4.01	4.48	4.80
Oil-price-pressure tone	359,607	0.23	0.91	-0.79	-0.39	0.09	0.71
Lagged bank controls							
Log assets, lagged	359,607	12.22	1.31	10.71	11.34	12.09	12.95
Equity/assets, lagged	359,607	11.53	6.21	7.80	8.90	10.36	12.46
Loans/assets, lagged	359,607	62.61	16.44	40.69	53.05	64.91	74.71
Deposits/assets, lagged	359,607	83.04	9.46	74.60	80.50	85.09	88.35
ROA, lagged	359,607	0.25	0.27	0.05	0.15	0.25	0.35

Notes: This table reports summary statistics for the matched bank-quarter sample used in the specifications. Bank outcomes, exposure shares, and balance-sheet ratios are reported in percentage points. WTI returns and WTI realized volatility are reported in percent. Log assets, standardized narrative indices, and oil-price-pressure tone are reported in their original standardized units. Deposit-weighted oil exposure is constructed from FDIC Summary of Deposits branch locations and county oil-and-gas payroll intensity from County Business Patterns. Narrative indices are standardized quarterly measures. Bank controls are lagged one quarter to match the empirical specification.

at least 10 percentage points. Appendix Table 9 reports the corresponding exposure distribution for all observations and for positive-exposure observations.

I also construct two robustness measures. The first sets Texas counties' oil intensity to zero while leaving Texas branch deposits in the denominator of deposit shares. This measure isolates exposure to non-Texas oil-intensive counties and tests whether the main results are driven entirely by Texas. The second classifies banks by whether their headquarters are located in major oil-producing states. This headquarters-state measure is simpler and more transparent, but coarser than the deposit-weighted exposure measure, so it is used as an alternative exposure definition rather than the baseline.

4 Empirical Design

The empirical design estimates whether banks with greater exposure to oil-intensive counties respond differently when particular oil narratives become salient. The baseline specification is a dynamic panel local projection (Jordà, 2005) with bank fixed effects and quarter fixed effects:

$$\begin{aligned} Y_{i,t+h} = \alpha_i + \delta_t + \beta_h \big(Narrative_t \times OilExposure_{i,t-1}\big) + \pi_h \, OilExposure_{i,t-1} \\ + \Lambda_h' \big(OilMarket_t \times OilExposure_{i,t-1}\big) + \Gamma_h' \, X_{i,t-1} + \varepsilon_{i,t+h}, \end{aligned} \tag{4}$$

where $Y_{i,t+h}$ is a bank outcome for bank i at horizon h, α_i denotes bank fixed effects, and δ_t denotes quarter fixed effects. The variable $Narrative_t$ is a quarterly oil narrative index. The variable $OilExposure_{i,t-1}$ is lagged oil exposure, constructed from FDIC Summary of Deposits branch locations and county oil-and-gas payroll intensity from County Business Patterns. The exposure variable is scaled so that coefficients correspond to a 10 percentage-point increase in oil exposure. The vector $OilMarket_t$ contains oil-market controls defined below, and $X_{i,t-1}$ contains lagged bank controls measured one quarter before the narrative quarter.

The coefficient of interest is β_h. It measures whether banks with greater oil exposure respond differently when a given oil narrative becomes more salient. The quarter fixed effects absorb the aggregate level of oil narratives, oil prices, macroeconomic conditions, monetary policy, and national banking shocks common to all banks in a given quarter. The bank fixed effects absorb persistent differences across institutions, including business model, geography, charter type, and long-run balance-sheet composition. Identification therefore comes from differential responses across banks with different degrees of oil-county exposure within the same quarter. The main effects of the quarterly narrative and oil-market variables are absorbed by the quarter fixed effects.

The baseline outcomes capture credit risk, balance-sheet growth, profitability, and capitalization. For flow variables such as provisions, charge-offs, loan growth, deposit growth, and return on assets, I estimate future cumulative local projections:

$$\begin{aligned}\sum_{j=1}^{h} Y_{i,t+j} = \alpha_i + \delta_t + \beta_h \big(Narrative_t \times OilExposure_{i,t-1}\big) + \pi_h \, OilExposure_{i,t-1} \\ + \Lambda_h' \big(OilMarket_t \times OilExposure_{i,t-1}\big) + \Gamma_h' \, X_{i,t-1} + \varepsilon_{i,t+h}^{(h)}.\end{aligned} \tag{5}$$

This cumulative specification excludes quarter t and summarizes the future balance-sheet response from quarter $t+1$ through quarter $t+h$. This timing avoids using outcomes that may partly occur before the quarterly narrative index is realized.

For stock variables such as nonperforming loans and equity/assets, I estimate changes from the pre-narrative balance sheet to the future horizon:

$$\begin{aligned}Y_{i,t+h} - Y_{i,t-1} = \alpha_i + \delta_t + \beta_h \big(Narrative_t \times OilExposure_{i,t-1}\big) + \pi_h \, OilExposure_{i,t-1} \\ + \Lambda_h' \big(OilMarket_t \times OilExposure_{i,t-1}\big) + \Gamma_h' \, X_{i,t-1} + \varepsilon_{i,t+h}^{(h)}.\end{aligned} \tag{6}$$

Using changes rather than future levels reduces the influence of persistence in bank balance-sheet ratios and anchors the stock response relative to the pre-narrative quarter.

The main specifications focus on two narrative measures: supply-disruption narratives and demand-salience narratives. Supply-disruption narratives capture coverage of outages, shortages, shutdowns, refinery constraints, sanctions, storms, and production interruptions. Demand-salience narratives capture coverage of demand, consumption, growth, China demand, global demand, and broader macroeconomic conditions. I estimate the two narrative interactions jointly:

$$Y_{i,t+h} = \alpha_i + \delta_t + \beta_h^S \left(Supply_t \times OilExposure_{i,t-1}\right) + \beta_h^D \left(Demand_t \times OilExposure_{i,t-1}\right) + \pi_h \, OilExposure_{i,t-1} + \Lambda_h' \left(OilMarket_t \times OilExposure_{i,t-1}\right) + \Gamma_h' \, X_{i,t-1} + \varepsilon_{i,t+h}. \tag{7}$$

The coefficients β_h^S and β_h^D distinguish oil narratives that operate as credit-risk signals from narratives that coincide with expansionary oil-region conditions. For cumulative flow outcomes, the dependent variable in equation (7) is replaced by $\sum_{j=1}^{h} Y_{i,t+j}$. For stock outcomes, it is replaced by $Y_{i,t+h} - Y_{i,t-1}$.

The baseline regression includes the following lagged bank controls: log assets, equity/assets, loans/assets, deposits/assets, and return on assets. These variables account for size, capitalization, balance-sheet structure, funding composition, and recent profitability. All bank controls are measured at quarter $t-1$.

A potential concern is that narrative measures may proxy for oil-price movements rather than distinct narrative information. To address this concern, the oil-market control vector is

$$OilMarket_t = \left(OilReturn_t, \, OilVolatility_t, \, OilTone_t\right)' . \tag{8}$$

In the regressions, each component of $OilMarket_t$ is interacted with oil exposure. These controls allow oil-exposed banks to respond differently to oil returns, realized oil-price volatility, and oil-price-pressure tone. The coefficients on supply-disruption and demand-salience narratives are therefore interpreted as differential bank responses to narrative salience beyond these oil-market controls.

Standard errors are two-way clustered by bank and quarter (Cameron et al., 2011; Bertrand et al., 2004). Clustering by bank allows for serial correlation in bank outcomes, while clustering by quarter allows for arbitrary cross-sectional dependence among banks within the same quarter. All regressions are estimated on the matched bank panel from 2001Q1 through 2024Q4.

5 Main Results

This section presents the main empirical results. I begin with aggregate oil-risk attention. The aggregate index is informative, but it produces a mixed balance-sheet pattern because it combines different oil narratives. I then decompose oil news into narrative types and show that the main contrast is between supply-disruption narratives and demand-salience narratives. Supply-disruption narratives predict credit-loss recognition and weaker balance-sheet growth among banks with greater exposure to oil-intensive counties. Demand-salience narratives predict stronger lending, stronger deposits, and higher profitability.

5.1 Aggregate Oil-Risk Attention

I first estimate the baseline panel local projection using aggregate oil-risk attention. This index measures the salience of risk-related oil language in newspaper coverage, without distinguishing the type of oil story. The specification interacts aggregate oil-risk attention with lagged oil exposure and includes bank fixed effects, quarter fixed effects, bank controls, and oil-market interactions.

Table 2 shows that aggregate oil-risk attention is informative, but not sufficient for a clear bank-fragility interpretation. Aggregate oil-risk attention does not predict higher provisions or charge-offs for more oil-exposed banks. Instead, the estimated responses are mixed: the NPL response is positive at the eight-quarter horizon, deposit growth rises at the four-quarter horizon, and profitability increases at both the four- and eight-quarter horizons. This pattern is difficult to interpret as a simple credit-risk channel.

The mixed aggregate result is useful as a diagnostic. It indicates that risk-related oil language contains information for banks exposed to oil-intensive counties, but it also suggests that the aggregate index pools oil stories with different banking implications. A single oil-risk index can combine supply disruptions, demand conditions, geopo-

Table 2: Aggregate Oil-Risk Attention and Bank Outcomes

	Horizon (quarters)	
Outcome	4	8
NPL ratio	0.01	0.04**
	(0.01)	(0.02)
Provisions/assets	−0.00	−0.02
	(0.01)	(0.02)
Charge-offs/assets	−0.00	−0.01
	(0.01)	(0.02)
Loan growth	0.41	0.58
	(0.41)	(0.73)
Deposit growth	0.45*	0.16
	(0.26)	(0.45)
ROA	0.04***	0.08***
	(0.01)	(0.03)
Equity/assets	−0.00	0.07
	(0.03)	(0.05)

Notes: This table reports panel local projection estimates using aggregate oil-risk attention interacted with lagged deposit-weighted oil exposure. Oil exposure is constructed from FDIC Summary of Deposits branch locations and county oil-and-gas payroll intensity from County Business Patterns. Coefficients are reported in percentage points and correspond to a one-standard-deviation increase in aggregate oil-risk attention interacted with a 10 percentage-point increase in deposit-weighted oil exposure. Flow outcomes are cumulative from quarter $t + 1$ through horizon h, excluding quarter t. Stock outcomes are changes from quarter $t - 1$ to quarter $t + h$. All specifications include bank fixed effects, quarter fixed effects, lagged bank controls, and oil-return, oil-volatility, and oil-tone interactions with deposit-weighted oil exposure. Standard errors, two-way clustered by bank and quarter, are in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

litical tension, inventory stress, and broader uncertainty. These stories need not have the same effect on exposed banks. Some may signal borrower stress and future credit losses. Others may coincide with stronger oil-sector income, stronger local demand, or expansion in oil-producing regions. The relevant question is therefore not whether oil news in general becomes bank risk, but which oil narratives do.

5.2 Narrative-Type Decomposition

I next separate oil narratives into supply-disruption and demand-salience measures. The regressions include both narrative interactions jointly, along with oil-return, oil-volatility, and oil-tone interactions with oil exposure. The coefficients therefore compare how banks with greater oil-county exposure respond to different narrative types after allowing exposed banks to respond differently to oil-price movements and oil-price-pressure tone.

Table 3 reports the main cumulative results. Supply-disruption narratives have a clear credit-risk effect. For banks with greater oil exposure, supply-disruption narratives predict higher cumulative loan-loss provisions and higher cumulative net charge-offs. The effects are statistically significant and grow over the horizon. By horizon eight, supply-disruption narratives raise cumulative provisions/assets by 0.11 percentage points and cumulative charge-offs/assets by 0.12 percentage points for a 10 percentage-point increase in deposit-weighted oil exposure.

Supply-disruption narratives also predict weaker balance-sheet growth. Cumulative loan growth falls by 0.77 percentage points at horizon four and by 1.81 percentage points at horizon eight. Cumulative deposit growth falls by 0.89 percentage points at horizon four and by 1.97 percentage points at horizon eight. Profitability also declines: cumulative ROA falls by 0.03 percentage points at horizon four and by 0.06 percentage points at horizon eight. These results indicate that supply-disruption narratives are associated not only with credit-loss recognition, but also with slower expansion and weaker profitability among oil-exposed banks.

Demand-salience narratives move in the opposite direction. They predict lower cu-

Table 3: Oil Narrative Types and Bank Outcomes

Outcome	Horizon (quarters)	Narrative channel: Supply disruption	Narrative channel: Demand salience
NPL ratio	4	0.06 * *	−0.06 * *
		(0.03)	(0.03)
	8	0.10 * **	−0.07 * *
		(0.03)	(0.03)
Provisions/assets	4	0.06 * **	−0.05 * **
		(0.01)	(0.01)
	8	0.11 * **	−0.12 * **
		(0.02)	(0.03)
Charge-offs/assets	4	0.06 * **	−0.05 * **
		(0.01)	(0.01)
	8	0.12 * **	−0.11 * **
		(0.03)	(0.03)
Loan growth	4	−0.77 * *	1.16 * **
		(0.31)	(0.39)
	8	−1.81 * **	2.39 * **
		(0.65)	(0.67)
Deposit growth	4	−0.89 * *	1.13 * **
		(0.38)	(0.35)
	8	−1.97 * **	1.72 * *
		(0.67)	(0.69)
ROA	4	−0.03 * *	0.07 * **
		(0.02)	(0.02)
	8	−0.06 * *	0.13 * **
		(0.03)	(0.04)
Equity/assets	4	0.03	0.01
		(0.05)	(0.05)
	8	0.12*	−0.00
		(0.06)	(0.08)

Notes: This table reports panel local projection estimates using lagged bank controls. Oil exposure is lagged deposit-weighted exposure to county oil-and-gas payroll intensity, constructed from FDIC Summary of Deposits branch locations and County Business Patterns. Coefficients are reported in percentage points and correspond to a one-standard-deviation increase in the narrative index interacted with a 10 percentage-point increase in deposit-weighted oil exposure. Flow outcomes are cumulative from quarter $t+1$ through horizon h, excluding quarter t. Stock outcomes are changes from quarter $t-1$ to quarter $t+h$. All specifications include bank fixed effects, quarter fixed effects, lagged bank controls, and oil-return, oil-volatility, and oil-tone interactions with deposit-weighted oil exposure. Standard errors, two-way clustered by bank and quarter, are in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

mulative provisions and charge-offs, stronger loan growth, stronger deposit growth, and higher profitability. At horizon eight, demand salience raises cumulative loan growth by 2.39 percentage points, deposit growth by 1.72 percentage points, and ROA by 0.13 percentage points, while reducing provisions and charge-offs. I interpret this index as demand salience rather than an adverse demand shock: it measures the prominence of oil-demand language, not only weak-demand news. In the bank data, demand salience is associated with expansionary conditions in oil-exposed banking markets.

The contrast between supply disruption and demand salience is the main empirical result. The same oil-news corpus contains narratives that have opposite implications for exposed banks. Supply-disruption narratives predict bank-risk outcomes, while demand-salience narratives appear to coincide with stronger local banking conditions. This explains why aggregate oil-risk attention produces mixed signs.

Because the narrative variables vary at the national quarterly frequency, I show in Section 6.4 that the supply-disruption pattern also survives a more aggregated exposure-bin-by-quarter design.

5.3 Dynamic Balance-Sheet Responses

Figure 3 plots the cumulative local projection coefficients for the two main narrative types. The figure focuses on cumulative provisions, cumulative charge-offs, cumulative loan growth, cumulative deposit growth, and cumulative return on assets. The shaded regions show 68 and 90 percent confidence intervals.

The dynamic responses reinforce the table results. Supply-disruption narratives generate a persistent increase in cumulative provisions and charge-offs among banks with greater oil exposure. The credit-loss response appears immediately and grows over the horizon. At the same time, supply-disruption narratives predict declining cumulative loan growth and deposit growth. This combination is consistent with a local bank-risk channel: more oil-exposed banks recognize more losses and expand more slowly after supply-disruption oil narratives intensify.

Demand-salience narratives generate the opposite pattern. Cumulative provisions and charge-offs decline, while loan growth, deposit growth, and profitability rise. This pattern is not consistent with demand salience acting as a bank-stress signal. Instead, it suggests that demand-related oil narratives often coincide with stronger oil-region banking conditions.

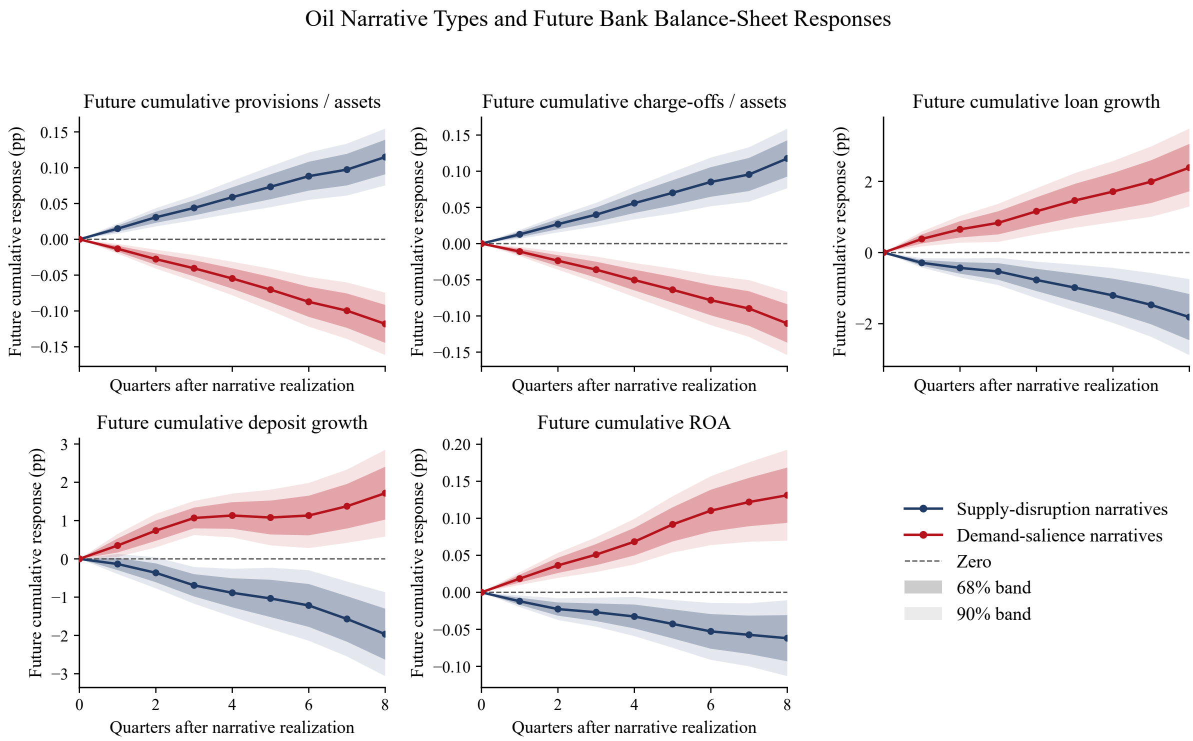

Figure 3: Oil Narrative Types and Bank Balance-Sheet Responses
Notes: The figure plots cumulative local projection coefficients for supply-disruption and demand-salience narratives interacted with lagged oil exposure. Bank controls are lagged one quarter. Flow outcomes are cumulative from quarter $t + 1$ through horizon h, excluding quarter t. Oil exposure is constructed from FDIC Summary of Deposits branch locations and county oil-and-gas payroll intensity from County Business Patterns. Coefficients are scaled to represent a one-standard-deviation increase in the narrative index interacted with a 10 percentage-point increase in oil exposure. All specifications include bank fixed effects, quarter fixed effects, lagged bank controls, and oil-return, oil-volatility, and oil-tone interactions with oil exposure. Shaded regions denote 68 and 90 percent confidence bands based on standard errors clustered by bank and quarter.

The main lesson is that oil narratives differentially predict bank outcomes through distinct exposure-based patterns. Supply-disruption narratives are associated with credit-risk realization and slower balance-sheet growth. Demand-salience narratives are associated with expansion. The banking effect of oil news depends on the narrative type, not simply on the volume of oil-risk language.

6 Robustness

This section examines whether the main narrative-type results are driven by oil-price movements, Texas counties, or the particular exposure definition used in the baseline analysis. The main finding is stable: supply-disruption narratives predict credit-loss recognition among banks with greater exposure to oil-intensive counties, while demand-salience narratives move in the opposite direction.

6.1 Oil-Price Controls

A central concern is that the narrative indices may proxy for oil-price movements rather than distinct information in newspaper coverage. This concern is especially important in the banking setting because banks exposed to oil-intensive counties may respond differently to oil-price increases, oil-price volatility, or news that implies upward price pressure.

The main narrative-type specifications therefore include oil-market controls interacted with lagged oil exposure. Specifically, the regressions include interactions between oil exposure and quarterly WTI returns, realized oil-price volatility, and oil-price-pressure tone. Quarter fixed effects absorb the aggregate oil-price environment common to all banks. The interacted controls allow more oil-exposed banks to respond differently to oil-price movements, oil-price volatility, and oil-price-pressure tone.

The supply-disruption and demand-salience coefficients are therefore not identified from a simple comparison of high-oil-price and low-oil-price quarters. They are estimated after allowing exposed banks to have different responses to oil returns, realized volatility, and oil-price-pressure tone. The persistence of the supply-disruption results under these controls suggests that the estimated credit-risk response reflects narrative content beyond contemporaneous oil-price movements.

6.2 Excluding Texas Counties from Exposure

Texas is the largest oil-producing state in the sample and contains many oil-intensive counties. A natural concern is that the baseline results could be driven entirely by Texas county exposure. I address this concern by constructing an alternative exposure measure that sets Texas counties' oil intensity to zero. Texas branch deposits remain in the denominator when computing deposit shares, so the alternative measure captures exposure to non-Texas oil-intensive counties.

Table 4 shows that the credit-risk channel is not driven only by Texas counties. When Texas counties are removed from the exposure construction, supply-disruption narratives continue to predict higher cumulative provisions and charge-offs. The effects remain statistically significant at both the four-quarter and eight-quarter horizons. This result is important because it shows that the main credit-loss finding is not simply a Texas-oil-bank result.

The balance-sheet-growth response is more sensitive to Texas exposure. The loan-growth contraction weakens substantially once Texas counties are removed from the exposure measure, while the deposit-growth response remains negative but smaller. This suggests that Texas exposure contributes importantly to the lending contraction in the full-sample results. The broader credit-risk result, however, remains intact: supply-disruption narratives still predict loss recognition and weaker profitability among banks exposed to non-Texas oil counties.

Demand salience remains strongly expansionary after excluding Texas counties from the exposure construction. It continues to predict lower cumulative provisions and charge-offs, stronger loan growth, stronger deposit growth, and higher profitability. This reinforces the central interpretation that demand-salience narratives do not operate as bank-stress signals. Instead, they coincide with stronger conditions in oil-exposed banking markets.

Table 4: Oil Narrative Types and Bank Outcomes: Excluding Texas Counties from Exposure

Outcome	Horizon (quarters)	Full exposure: Supply disruption	Full exposure: Demand salience	Excluding Texas counties: Supply disruption	Excluding Texas counties: Demand salience
Provisions/assets	4	0.06***	−0.05***	0.06***	−0.06***
		(0.01)	(0.01)	(0.01)	(0.01)
	8	0.11***	−0.12***	0.11***	−0.12***
		(0.02)	(0.03)	(0.02)	(0.03)
Charge-offs/assets	4	0.06***	−0.05***	0.05***	−0.05***
		(0.01)	(0.01)	(0.01)	(0.01)
	8	0.12***	−0.11***	0.11***	−0.11***
		(0.03)	(0.03)	(0.03)	(0.03)
Loan growth	4	−0.77**	1.16***	0.01	1.02**
		(0.31)	(0.39)	(0.36)	(0.47)
	8	−1.81***	2.39***	−0.21	1.90***
		(0.65)	(0.67)	(0.82)	(0.73)
Deposit growth	4	−0.89**	1.13***	−0.72*	0.99***
		(0.38)	(0.35)	(0.39)	(0.36)
	8	−1.97***	1.72**	−1.49*	1.54*
		(0.67)	(0.69)	(0.76)	(0.79)
ROA	4	−0.03**	0.07***	−0.04**	0.09***
		(0.02)	(0.02)	(0.02)	(0.02)
	8	−0.06**	0.13***	−0.09**	0.19***
		(0.03)	(0.04)	(0.04)	(0.04)

Notes: This table compares cumulative local projection estimates using the full deposit-weighted oil exposure measure and an alternative measure that sets Texas counties' oil intensity to zero. Branch deposits in Texas remain in the denominator when constructing deposit shares, so the alternative exposure captures banks' deposit-weighted exposure to non-Texas oil counties. Coefficients are reported in percentage points and correspond to a one-standard-deviation increase in the narrative index interacted with a 10 percentage-point increase in deposit-weighted oil exposure. Flow outcomes are cumulative from quarter $t+1$ through horizon h, excluding quarter t. All specifications include bank fixed effects, quarter fixed effects, lagged bank controls, and oil-return, oil-volatility, and oil-tone interactions with the relevant deposit-weighted oil exposure measure. Standard errors, two-way clustered by bank and quarter, are in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

6.3 Headquarters-State Exposure

As an alternative exposure definition, I also estimate the dynamic responses using a simpler headquarters-state measure. This measure classifies a bank as oil-state exposed if its headquarters is located in a major oil-producing state. The headquarters-state measure is transparent and available for the full bank panel, but it is coarser than deposit-weighted exposure because it assigns each bank to a single state and does not use the geographic distribution of branch deposits.

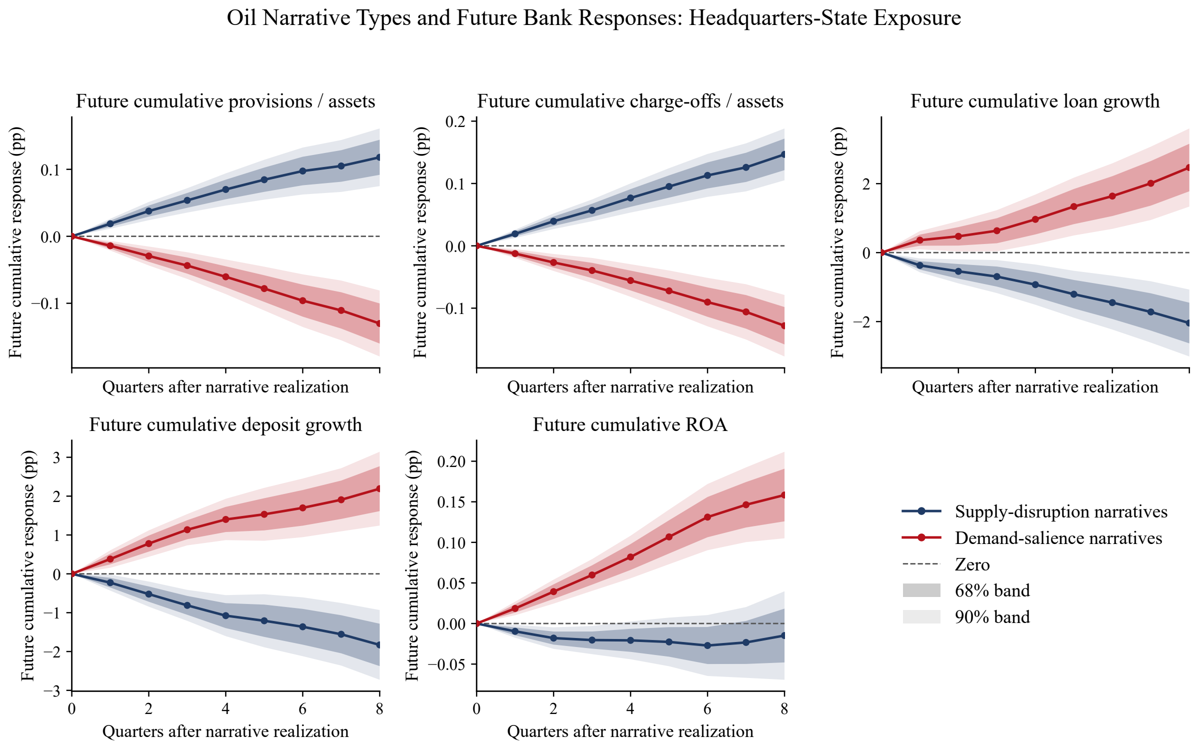

Figure 4: Oil Narrative Types and Bank Balance-Sheet Responses: Headquarters-State Exposure

Notes: The figure plots cumulative local projection coefficients for supply-disruption and demand-salience narratives interacted with a headquarters-state oil-exposure indicator. Bank controls are lagged one quarter. Flow outcomes are cumulative from quarter $t+1$ through horizon h, excluding quarter t. All specifications include bank fixed effects, quarter fixed effects, lagged bank controls, and oil-return, oil-volatility, and oil-tone interactions with the headquarters-state exposure indicator. Shaded regions denote 68 and 90 percent confidence bands based on standard errors clustered by bank and quarter.

Figure 4 reports the corresponding dynamic responses using the headquarters-state exposure definition. The qualitative pattern is similar to the baseline results. Supply-disruption narratives predict higher provisions and charge-offs and weaker balance-sheet growth for oil-state banks, while demand-salience narratives move in the opposite direction. The similarity between the headquarters-state results and the deposit-weighted results suggests that the main findings are not an artifact of a single exposure construction.

6.4 Additional Robustness Tests

I report three further robustness tests in Appendix B. First, the main results are not driven by the two largest oil-market episodes in the sample. Appendix Table 11 re-estimates the narrative-type specification after excluding the 2014–2016 oil-price collapse and the 2020 COVID oil-market episode, and the supply-disruption and demand-salience patterns are unchanged. Second, the results hold across bank size. Appendix Table 12 splits the panel by bank size and shows that both small and large banks display the supply-disruption credit-risk response and the demand-salience expansion. Third, the results do not depend on the large number of bank-quarter observations. Because the narrative indices vary at the national quarterly frequency, Appendix Table 13 collapses the panel to exposure-bin-by-quarter cells and estimates the timing-clean local projections with Driscoll-Kraay standard errors (**?**). The supply-disruption pattern remains: more exposed bins show higher nonperforming loans, provisions, and charge-offs, along with weaker deposit growth. Appendix Table 10 reports sample sizes for the main estimates.

7 Conclusion

Oil narratives do not have one banking effect. This paper shows that the balance-sheet consequences of oil news depend on the type of oil story that becomes salient. Using newspaper-based oil narrative measures, quarterly U.S. bank balance-sheet data from 2001 to 2024, and exposure to oil-intensive counties, I estimate dynamic panel local projections that compare banks with different degrees of local oil exposure within the same quarter.

The main finding is that supply-disruption narratives is associated with bank risk.

When newspaper coverage emphasizes outages, shortages, production interruptions, refinery constraints, sanctions, storms, and related supply disruptions, more oil-exposed banks experience higher cumulative provisions and charge-offs, lower loan growth, lower deposit growth, and weaker short-run profitability. These effects are not explained by oil returns, realized oil-price volatility, or oil-price-pressure tone. They also survive a stricter exposure test that removes Texas counties from the construction of oil exposure, although the lending contraction is weaker once Texas exposure is removed.

Demand-salience narratives have a different effect. They are associated with stronger lending, stronger deposit growth, and higher profitability among more oil-exposed banks, while provisions and charge-offs decline. This does not mean that all demand-related oil news is favorable. Rather, the demand-salience index captures the prominence of oil-demand language, including coverage of global growth, consumption, China demand, demand recovery, and broader macroeconomic conditions. In the bank data, this narrative type is associated with expansionary conditions in oil-exposed banking markets.

The results explain why aggregate oil-risk attention produces a mixed balance-sheet pattern. A single oil-risk index pools narratives that have different implications for exposed banks. Some oil stories signal borrower stress and future credit losses. Others coincide with stronger regional activity. The banking effect of oil news therefore depends not simply on whether oil markets are receiving attention, but on which narrative is driving that attention.

The broader implication is that narrative measurement matters for financial-stability analysis. Banks operate in information environments where news about disruptions, demand, uncertainty, and market stress shapes expectations about borrowers and local conditions. Treating oil news as a single aggregate signal can obscure economically important heterogeneity. Distinguishing narrative types reveals when oil-market information becomes bank risk and when it instead reflects expansion in oil-exposed regions.

Several extensions are natural. Future work could examine whether oil narratives affect bank stock returns, loan pricing, credit-line drawdowns, or supervisory risk measures. Another useful extension would separate energy-loan exposure from deposit-location exposure, allowing a sharper distinction between borrower-credit channels and local funding channels. The evidence here provides a first step: oil narratives signal to banks, but not all oil narratives signal in the same way.

References

Abeyie, Gabriel, "Not All Oil Narrative Shocks Are Alike: The Macroeconomic Effects of News-Driven Sentiment and Uncertainty Shocks," *Available at SSRN 6532699*, 2025.

Agoraki, Maria-Eleni K., Nektarios Aslanidis, and Georgios P. Kouretas, "U.S. Banks' Lending, Financial Stability, and Text-Based Sentiment Analysis," *Journal of Economic Behavior & Organization*, 2022, *197*, 73–90.

Bae, Junsung, Allen N. Berger, Hyun-Soo Choi, and Hugh H. Kim, "Bank Sentiment, Loan Loss Provisioning, and Lending," 2024. Working paper.

Baker, Scott R., Nicholas Bloom, and Steven J. Davis, "Measuring Economic Policy Uncertainty," *Quarterly Journal of Economics*, 2016, *131* (4), 1593–1636.

Baumeister, Christiane and James D. Hamilton, "Structural Interpretation of Vector Autoregressions with Incomplete Identification: Revisiting the Role of Oil Supply and Demand Shocks," *American Economic Review*, 2019, *109* (5), 1873–1910.

Becker, Bo, "Geographical Segmentation of US Capital Markets," *Journal of Financial Economics*, 2007, *85* (1), 151–178.

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan, "How Much Should We Trust Differences-in-Differences Estimates?," *Quarterly Journal of Economics*, 2004, *119* (1), 249–275.

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller, "Robust Inference with Multiway Clustering," *Journal of Business and Economic Statistics*, 2011, *29* (2), 238–249.

Cortés, Kristle Romero and Philip E. Strahan, "Tracing Out Capital Flows: How Financially Integrated Banks Respond to Natural Disasters," *Journal of Financial Economics*, 2017, *125* (1), 182–199.

Dell'Atti, Stefano, Andrea Paltrinieri, Caterina Di Tommaso, and Grazia Onorato, "Assessment of Banking Risk in the Context of the Oil and Gas Bubbles," *Energy Economics*, 2025, *147.*

Driscoll, John C and Aart C Kraay, "Consistent covariance matrix estimation with spatially dependent panel data," *Review of economics and statistics*, 1998, *80* (4), 549–560.

García, Diego, "Sentiment during Recessions," *Journal of Finance*, 2013, *68* (3), 1267–1300.

Gentzkow, Matthew, Bryan Kelly, and Matt Taddy, "Text as Data," *Journal of Economic Literature*, 2019, *57* (3), 535–574.

Gilje, Erik P., "Does Local Access to Finance Matter? Evidence from U.S. Oil and Natural Gas Shale Booms," *Management Science*, 2019, *65* (1), 1–18.

_ , Elena Loutskina, and Philip E. Strahan, "Exporting Liquidity: Branch Banking and Financial Integration," *Journal of Finance*, 2016, *71* (3), 1159–1184.

Hamilton, James D., "Oil and the Macroeconomy since World War II," *Journal of Political Economy*, 1983, *91* (2), 228–248.

Jordà, Òscar, "Estimation and Inference of Impulse Responses by Local Projections," *American Economic Review*, 2005, *95* (1), 161–182.

Känzig, Diego R., "The Macroeconomic Effects of Oil Supply News: Evidence from OPEC Announcements," *American Economic Review*, 2021, *111* (4), 1092–1125.

Kilian, Lutz, "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market," *American Economic Review*, 2009, *99* (3), 1053–1069.

Landier, Augustin, David Sraer, and David Thesmar, "Banking Integration and House Price Co-movement," *Journal of Financial Economics*, 2017, *125* (1), 1–25.

Loughran, Tim and Bill McDonald, "When Is a Liability Not a Liability? Textual Analysis, Dictionaries, and 10-Ks," *Journal of Finance*, 2011, *66* (1), 35–65.

Manela, Asaf and Alan Moreira, "News Implied Volatility and Disaster Concerns," *Journal of Financial Economics*, 2017, *123* (1), 137–162.

Morgan, Donald P., Bertrand Rime, and Philip E. Strahan, "Bank Integration and State Business Cycles," *Quarterly Journal of Economics*, 2004, *119* (4), 1555–1584.

Ng, Jeffrey, Walid Saffar, and Janus Jian Zhang, "Policy Uncertainty and Loan Loss Provisions in the Banking Industry," *Review of Accounting Studies*, 2020, *25* (2), 726–777.

Pacelli, Vincenzo, Caterina Di Tommaso, Matteo Foglia, and Maria Melania Povia, "Spillover Effects between Energy Uncertainty and Financial Risk in the Eurozone Banking Sector," *Energy Economics*, 2025, *141.*

Shiller, Robert J., "Narrative Economics," *American Economic Review*, 2017, *107* (4), 967–1004.

Tetlock, Paul C., "Giving Content to Investor Sentiment: The Role of Media in the Stock Market," *Journal of Finance*, 2007, *62* (3), 1139–1168.

Wang, Teng, "Local Banks and the Effects of Oil Price Shocks," *Journal of Banking and Finance*, 2021, *125*, 106069.

A Text-Measure Construction

This appendix summarizes the construction of the newspaper-based oil narrative measures. The empirical analysis uses quarterly narrative indices constructed from monthly oil-news coverage. The text measures are designed to capture the salience of different oil-market stories rather than to identify structural oil shocks.

I begin from oil-related newspaper articles and construct monthly counts of narrative-specific language. Supply-disruption narratives count terms associated with outages, shortages, shutdowns, production interruptions, refinery constraints, sanctions, storms, and related disruptions. Demand-salience narratives count terms associated with oil demand, consumption, global growth, China demand, demand recovery, and broader macroeconomic conditions. Aggregate oil-risk attention counts risk-related oil language. Oil-price-pressure tone measures the direction of price-pressure language in oil-news coverage and is used as a control in the bank regressions.

The baseline indices use log-transformed monthly counts and are standardized over the monthly sample. I aggregate the standardized monthly measures to the quarterly frequency by averaging within quarter. The resulting variables enter the bank regressions as standardized quarterly narrative measures. Thus, the reported coefficients correspond to a one-standard-deviation increase in narrative salience interacted with a 10 percentage-point increase in deposit-weighted oil exposure.

The demand-salience index should not be interpreted as a negative oil-demand shock. It is a salience measure: it captures periods when oil-demand language becomes more prominent in coverage. This language can reflect demand weakness, but it can also reflect global growth, demand recovery, and strong consumption. For this reason, the paper interprets positive demand-salience coefficients as evidence that demand-related oil coverage often coincides with expansionary conditions in oil-exposed banking markets, not as evidence that adverse demand shocks improve bank outcomes.

Appendix Table 5 summarizes the construction and empirical use of the main text and oil-market measures.

Table 5: Construction of Oil Narrative Measures

Measure	Concept	Monthly construction	Quarterly use
Aggregate oil-risk attention	Risk-related language in oil-news coverage.	Count risk-related oil terms; transform and standardize over the monthly sample.	Averaged within quarter and interacted with lagged deposit-weighted oil exposure.
Supply-disruption narratives	Oil stories about outages, shortages, shutdowns, refinery constraints, sanctions, storms, and production interruptions.	Count supply-disruption terms in oil-news coverage; log-transform counts and standardize.	Main narrative index; estimated jointly with demand salience.
Demand-salience narratives	Oil stories about demand, consumption, global growth, China demand, demand recovery, and macroeconomic conditions.	Count demand-related terms in oil-news coverage; log-transform counts and standardize.	Main narrative index; interpreted as demand salience, not a structural adverse demand shock.
Oil-price-pressure tone	Direction of price-pressure language in oil-news coverage.	Construct tone from oil-news language indicating upward versus downward oil-price pressure; standardize.	Included as an oil-market control interacted with deposit-weighted oil exposure.
Oil returns	Realized oil-price movement.	Quarterly WTI return.	Included as an oil-market control interacted with deposit-weighted oil exposure.
Oil volatility	Realized oil-market volatility.	Quarterly realized WTI volatility.	Included as an oil-market control interacted with deposit-weighted oil exposure.

Notes: This table summarizes the text and oil-market measures used in the empirical analysis. The narrative measures are constructed monthly from oil-news coverage and then averaged to the quarterly frequency. The narrative indices are standardized, so coefficient magnitudes correspond to a one-standard-deviation increase in narrative salience. Oil returns, realized oil volatility, and oil-price-pressure tone are included to distinguish narrative salience from contemporaneous oil-price movements and price-pressure language.

As a validation exercise, Appendix Table 6 reports the behavior of the narrative indices around major oil-market episodes. Appendix Table 7 lists the highest months for each index. These diagnostics are not used for estimation; they show whether the text measures spike during recognizable oil-news episodes and whether the supply-disruption and demand-salience measures capture distinct types of oil coverage.

Table 6: Oil Narrative Indices Around Major Oil-Market Episodes

Episode	Window	Expected narrative	Supply mean	Supply max	Demand mean	Demand max
Gulf War	1990-08 to 1991-03	Supply/geopolitical disruption	1.68	3.36	1.77	3.17
Katrina/Rita	2005-08 to 2005-10	Domestic supply disruption	2.43	3.47	1.93	2.22
Oil price peak	2008-05 to 2008-07	Demand and market pressure	1.05	1.23	2.32	2.53
Libya civil war	2011-02 to 2011-09	Supply/geopolitical disruption	-0.19	0.79	0.61	1.36
Oil-price collapse / OPEC strategy	2014-11 to 2016-02	Supply/demand re-balancing	0.64	1.41	1.12	2.46
COVID-19 oil shock	2020-03 to 2020-06	Demand collapse / market stress	0.90	1.63	1.20	1.79
Russia–Ukraine invasion	2022-02 to 2022-06	Supply/geopolitical disruption	0.56	1.10	0.79	1.63

Notes: This table reports average and maximum standardized narrative-index values during selected oil-market episodes. Supply-disruption narratives capture language related to outages, shortages, refinery constraints, sanctions, storms, and production interruptions. Demand-salience narratives capture language related to oil demand, consumption, growth, China demand, demand recovery, and broader macroeconomic conditions. The table is intended as a validation exercise: the indices should rise during episodes in which the corresponding narrative was prominent.

B Additional Robustness Results

This appendix also reports additional diagnostics and robustness checks for the local projection estimates. Appendix Table 10 reports sample sizes for the main estimates. Appendix Table 11 shows that the main results are not driven by the 2014–2016 oil-price collapse or the 2020 COVID/oil-market episode. Appendix Table 12 reports heterogeneity by bank size. The main supply-disruption and demand-salience patterns appear for both small and large banks. Because the narrative variables vary at the national quarterly frequency, Appendix Table 13 also reports estimates after collaps-

Table 7: Highest Narrative-Salience Months

Supply-disruption index			Demand-salience index		
Month	Supply	Demand	Month	Demand	Supply
2005-09	3.47	2.09	1990-08	3.17	3.36
1990-08	3.36	3.17	2004-08	2.74	1.55
2005-08	2.36	2.22	2004-10	2.54	1.30
2008-09	2.14	1.45	2008-06	2.53	1.23
2022-10	2.12	1.39	2015-01	2.46	1.17
2000-03	2.12	1.32	2008-07	2.40	0.94
2006-10	2.05	1.16	1990-09	2.39	2.03
1990-09	2.03	2.39	2008-10	2.37	1.26
2001-11	1.98	1.19	2014-12	2.29	1.37
1988-04	1.90	0.57	2005-08	2.22	2.36
2000-09	1.89	1.62	2016-01	2.19	0.95
1988-11	1.89	0.49	2008-12	2.10	1.50
1991-03	1.89	0.89	2005-09	2.09	3.47
2004-09	1.82	1.96	2008-05	2.02	0.98
1991-01	1.63	1.83	2004-05	1.99	1.58

Notes: This table lists the highest months for the supply-disruption and demand-salience narrative indices. The table provides a transparent diagnostic for whether the largest index values correspond to recognizable oil-news episodes.

ing the data to exposure-bin-by-quarter cells; the supply-disruption credit-risk pattern remains intact.

Table 8: Support for Deposit-Weighted Oil Exposure

Exposure threshold	Observations/Banks	Share (%)	Unique banks
Bank-quarter observations			
> 0	114,827	31.58	2,053
$\geq$ 0.01 pp	92,846	25.54	1,528
$\geq$ 0.1 pp	73,423	20.19	1,215
$\geq$ 0.5 pp	56,397	15.51	942
$\geq$ 1 pp	45,626	12.55	804
$\geq$ 2.5 pp	31,623	8.70	615
$\geq$ 5 pp	19,067	5.24	440
$\geq$ 10 pp	9,427	2.59	238
$\geq$ 15 pp	5,584	1.54	172
Banks ever reaching threshold			
> 0	2,053	48.35	2,053
$\geq$ 0.01 pp	1,528	35.99	1,528
$\geq$ 0.1 pp	1,215	28.62	1,215
$\geq$ 0.5 pp	942	22.19	942
$\geq$ 1 pp	804	18.94	804
$\geq$ 2.5 pp	615	14.48	615
$\geq$ 5 pp	440	10.36	440
$\geq$ 10 pp	238	5.61	238
$\geq$ 15 pp	172	4.05	172

Notes: This table reports support for the lagged deposit-weighted oil exposure measure. Exposure is measured in percentage points and is constructed from FDIC Summary of Deposits branch locations and county oil-and-gas payroll intensity from County Business Patterns. The upper panel reports bank-quarter observations above each threshold. The lower panel reports the number of banks that ever reach each threshold during the sample. The baseline estimates are scaled to a 10 percentage-point increase in exposure, so the table documents the empirical support for that magnitude.

Table 9: Distribution of Deposit-Weighted Oil Exposure

Sample	N	Mean	SD	P25	Median	P75	P90	P95
All bank-quarters	363,585	0.93	3.68	0.00	0.00	0.01	1.86	5.26
Positive-exposure bank-quarters	114,827	2.94	6.09	0.03	0.46	2.95	8.49	14.73
Bank-level mean exposure	4,246	0.90	3.10	0.00	0.00	0.06	2.23	5.55
Bank-level maximum exposure	4,246	1.90	6.08	0.00	0.00	0.22	5.35	11.28
Bank-level mean exposure, ever positive	2,053	1.86	4.26	0.00	0.07	1.54	5.68	9.94
Bank-level maximum exposure, ever positive	2,053	3.94	8.27	0.01	0.25	3.87	11.91	20.81

Notes: This table reports the distribution of lagged deposit-weighted oil exposure in percentage points. Positive-exposure bank-quarters restrict the sample to observations with strictly positive exposure. Bank-level mean exposure averages exposure across each bank's observed quarters, while bank-level maximum exposure records the largest exposure observed for each bank.

Table 10: Regression Diagnostics for Local Projection Estimates

Outcome	Horizon	Observations	Banks	Quarters
NPL ratio	4	341,050	4,203	88
NPL ratio	8	324,486	4,190	84
Provisions/assets	4	342,164	4,230	88
Provisions/assets	8	324,881	4,217	84
Charge-offs/assets	4	342,164	4,230	88
Charge-offs/assets	8	324,881	4,217	84
Loan growth	4	342,917	4,231	88
Loan growth	8	326,256	4,218	84
Deposit growth	4	342,917	4,231	88
Deposit growth	8	326,256	4,218	84
ROA	4	342,164	4,230	88
ROA	8	324,881	4,217	84
Equity/assets	4	342,367	4,230	88
Equity/assets	8	325,374	4,217	84

Notes: This table reports the number of observations, banks, and quarters used in the local projection estimates reported in Table 3. Flow outcomes are cumulative from quarter $t+1$ through horizon h; stock outcomes are changes from quarter $t-1$ to quarter $t+h$.

Table 11: Appendix Robustness: Episode Exclusions

	Full sample		Excl. 2014–2016		Excl. 2020	
Outcome	Supply	Demand	Supply	Demand	Supply	Demand
Provisions/assets	0.11***	−0.12***	0.10***	−0.14***	0.14***	−0.13***
	(0.02)	(0.03)	(0.03)	(0.03)	(0.02)	(0.03)
Charge-offs/assets	0.12***	−0.11***	0.12***	−0.14***	0.15***	−0.13***
	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)
Loan growth	−1.81***	2.39***	−1.79**	3.16***	−2.25***	2.94***
	(0.65)	(0.67)	(0.84)	(0.85)	(0.73)	(0.77)
Deposit growth	−1.97***	1.72**	−1.49**	1.60*	−2.52***	2.41***
	(0.67)	(0.69)	(0.74)	(0.85)	(0.69)	(0.71)
ROA	−0.06**	0.13***	−0.05	0.18***	−0.10***	0.15***
	(0.03)	(0.04)	(0.04)	(0.05)	(0.03)	(0.04)

Notes: This table reports horizon-eight cumulative local projection estimates under episode exclusions. The 2014–2016 exclusion removes quarters whose narrative quarter or future outcome window overlaps 2014Q1–2016Q4. The 2020 exclusion removes quarters whose narrative quarter or future outcome window overlaps 2020Q1–2020Q4. Coefficients are reported in percentage points and correspond to a one-standard-deviation increase in the narrative index interacted with a 10 percentage-point increase in lagged deposit-weighted oil exposure. Flow outcomes are cumulative from quarter $t+1$ through horizon h, excluding quarter t. All specifications include bank fixed effects, quarter fixed effects, lagged bank controls, and oil-return, oil-volatility, and oil-tone interactions with deposit-weighted oil exposure. Standard errors are two-way clustered by bank and quarter. ***, **, and * denote significance at the 1%, 5%, and 10% levels.

Table 12: Appendix Heterogeneity: Bank Size

	Small banks		Large banks	
Outcome	Supply	Demand	Supply	Demand
Provisions/assets	0.09***	−0.09***	0.15***	−0.15***
	(0.02)	(0.02)	(0.03)	(0.03)
Charge-offs/assets	0.09***	−0.09***	0.16***	−0.13***
	(0.02)	(0.02)	(0.03)	(0.03)
Loan growth	−1.92**	2.86***	−1.45**	1.27*
	(0.83)	(0.79)	(0.61)	(0.68)
Deposit growth	−1.39**	1.61***	−2.83***	1.59*
	(0.66)	(0.61)	(0.76)	(0.90)
ROA	−0.08**	0.14***	−0.04	0.12***
	(0.03)	(0.04)	(0.04)	(0.04)

Notes: This table reports horizon-eight cumulative local projection estimates by bank size. Small and large banks are defined using the median of each bank's first observed log assets. Coefficients are reported in percentage points and correspond to a one-standard-deviation increase in the narrative index interacted with a 10 percentage-point increase in lagged deposit-weighted oil exposure. Flow outcomes are cumulative from quarter $t+1$ through horizon h, excluding quarter t. All specifications include bank fixed effects, quarter fixed effects, lagged bank controls, and oil-return, oil-volatility, and oil-tone interactions with deposit-weighted oil exposure. Standard errors are two-way clustered by bank and quarter. ***, **, and * denote significance at the 1%, 5%, and 10% levels.

Table 13: Robustness: Exposure-Bin by Quarter Estimates

Outcome	Horizon (quarters)	Narrative channel: Supply disruption	Narrative channel: Demand salience
NPL ratio	4	0.06	−0.05
		(0.04)	(0.04)
	8	0.08***	−0.04
		(0.03)	(0.03)
Provisions/assets	4	0.05**	−0.01
		(0.02)	(0.02)
	8	0.12***	−0.06
		(0.04)	(0.04)
Charge-offs/assets	4	0.05***	−0.02
		(0.02)	(0.02)
	8	0.13***	−0.06
		(0.03)	(0.03)
Loan growth	4	−0.30	0.06
		(0.30)	(0.33)
	8	−1.04*	0.60
		(0.54)	(0.70)
Deposit growth	4	−0.78**	0.34
		(0.37)	(0.52)
	8	−1.48***	0.12
		(0.56)	(0.84)
ROA	4	−0.02	0.02
		(0.02)	(0.02)
	8	−0.06	0.04
		(0.05)	(0.05)
Equity/assets	4	−0.00	0.12*
		(0.05)	(0.07)
	8	0.02	0.18*
		(0.07)	(0.10)

Notes: This table reports local projection estimates after collapsing the bank panel to exposure-bin by quarter cells. Exposure bins are zero exposure, low positive exposure, medium exposure, and high exposure, based on lagged deposit-weighted oil exposure. The regressions include exposure-bin fixed effects, quarter fixed effects, lagged bin-average bank controls, and oil-return, oil-volatility, and oil-tone interactions with bin-average oil exposure. Cells are weighted by the number of bank-quarter observations in each exposure-bin by quarter cell. Standard errors are Driscoll and Kraay (1998) standard errors with a Bartlett kernel and bandwidth four. Coefficients are reported in percentage points and correspond to a one-standard-deviation increase in the narrative index interacted with a 10 percentage-point increase in deposit-weighted oil exposure. Flow outcomes are cumulative from quarter $t+1$ through horizon h; stock outcomes are changes from quarter $t-1$ to quarter $t+h$.