Transmission of Climate Risk in Corporate Loans*

Natalia Madrid† Max Yang‡

September 26, 2026

Abstract

Local natural disasters are geographically concentrated, but corporate credit markets are integrated through national financial intermediaries, raising the question of whether natural disasters propagate through lending networks. In August 2017, Hurricane Harvey caused widespread damage to counties in Texas and Louisiana. Using U.S. syndicated loan data, we compare lenders with different pre-storm portfolio concentration in the damaged counties lending to the same unaffected borrower. On the intensive margin, lenders one standard deviation more exposed to Harvey cut credit to unaffected borrowers by 4.8 percent of loan volume, with no adjustment in spreads or maturities. The effect is also not explained by distance to the storm, and it is unchanged among borrowers located more than 1,000 miles from the impact zone. Instead, the borrower's hurricane risk predicts the credit supply contraction, which we measure using the historical frequency of hurricanes at the borrower's location from NOAA data. Credit falls for borrowers in hurricane-prone counties, and exposed lenders move about two percentage points of their portfolios out of hurricane-risky counties. On the extensive margin, the most exposed lenders are more likely to end lending relationships with hurricane-risky borrowers. The evidence is consistent with a highly visible disaster leading exposed lenders to move away from climate-risky borrowers throughout their portfolios, a reallocation of climate risk in quantities.

JEL Classification: G21, G23, Q51, Q54

Keywords: Banks, Credit supply, Natural disasters, Climate risk, Syndicated loans

*We thank Judson Boomhower, Allan Timmermann, Mark Jacobsen, Carlos Avenancio-León, Juan Herreño, and Michael Reher for their invaluable advice and support. We also thank Robert Engle, Willian Mullins, and participants at the NBER Climate Finance Conference, Seminars at Banco de Mexico, and UC San Diego internal seminars for helpful comments. All errors are our own.

†Economics Department, University of California San Diego. Email: namadrid@ucsd.edu.

‡Rady School of Management, University of California San Diego. Email: ygnmax@ucsd.edu.

1 Introduction

Local climate disasters are increasingly severe and hard to fully insure or diversify. Hurricanes, floods, and wildfires destroy physical capital, disrupt production, and depress local demand, generating large and persistent economic losses (e.g., Garmaise and Moskowitz, 2009; Billings et al., 2022; Collier et al., 2024; Gallagher and Hartley, 2025). Over the last decade alone, the United States has experienced nearly 190 billion-dollar weather and climate disasters, with total damages exceeding $1.4 trillion, and these costs continue to rise (NOAA, 2025). The difficulty of insuring these losses comes from their structure. Insurance and diversification work by pooling many independent risks, so that a loss in one place is offset by premiums or gains elsewhere. Disaster losses break this logic in two ways. First, they are correlated in space. A hurricane damages every home, firm, and bank branch in the same set of counties at once, so an insurer or a lender doing business in the region cannot offset one loss against another nearby. Second, they are concentrated in time. Most years bring little damage and a single year brings an enormous amount, so losses may not be smoothed across years either. In practice, insurance against these losses is incomplete, and uninsured borrowers bear substantial financial distress when disasters strike (Billings et al., 2022).

Even when disasters are geographically concentrated, modern credit markets are integrated through national financial intermediaries, raising the question of whether natural disasters propagate through lending networks. A lender that funds firms in a disaster region also funds firms everywhere else, so a loss that cannot be pooled where it occurs is absorbed on a balance sheet that lends across the country. Propagation, however, is not automatic. A disaster also raises credit demand in the damaged region, because firms and households borrow to rebuild. If external funding were costless, a lender facing that demand would simply raise new deposits or wholesale funding, expand its balance sheet, and leave every other borrower untouched. But external finance is costly, so lenders meet local demand surges in part by reallocating funds within the balance sheet they already have, drawing credit away from markets far from the shock (Cortés and Strahan, 2017). And if the lenders absorbing the loss also reassess the climate risk they hold, they may withdraw specifically from climate-risky borrowers far from the disaster, turning a concentrated regional shock into an economy-wide reallocation of credit. Regulators share this concern, that idiosyncratic risks become systemic

risks. In June 2017, the Financial Stability Board's Task Force on Climate-related Financial Disclosures (TCFD) issued the first framework asking financial institutions to measure and disclose their exposure to climate risk (Task Force on Climate-related Financial Disclosures, 2017), and later that year, a group of central banks created the Network for Greening the Financial System (NGFS) to understand how climate risks affect the financial system and the economy (Network for Greening the Financial System, 2018). This propagation is what climate stress tests are built to capture, and assessing it requires knowing how these shocks actually travel through the financial system.

Most of the literature on climate change and financial markets has concentrated on long-lived assets, such as real estate or equities, whose prices should capitalize future long-term damages. The logic is that an asset held for decades must embed the decades of physical risk that come with it, so its price is the natural place to read off what markets believe about climate change. This literature has established that the discount is real and that it responds to information. Coastal properties exposed to sea-level rise trade at a discount relative to otherwise identical inland properties (Bernstein et al., 2019), the discount widens with the salience of flood risk and with the buyer's own beliefs, and lenders adjust the mortgages written against these properties by securitizing them rather than holding them (Ouazad and Kahn, 2022). Municipal debt issued by exposed jurisdictions carries a yield premium (Goldsmith-Pinkham et al., 2023), though the capitalization is not universal, since Murfin and Spiegel (2020) find precisely estimated null effects of sea-level-rise exposure on residential prices. Housing markets have even been used to recover the term structure of discount rates relevant for valuing climate abatement itself (Giglio et al., 2015, 2021). The common feature is that the horizon of the claim is long enough for a gradual, forward-looking risk to matter for its valuation today.

Much less is known about short-term assets, such as corporate loans, even though these are the claims intermediaries can reprice and resize fastest. A syndicated credit is renegotiated, amended, and refinanced continuously, so a lender that revises its view of climate risk can act on that view within quarters rather than waiting decades for a maturity. What the literature on corporate credit and disasters has mostly examined is the lending response inside the damaged regions themselves. Damaged firms and households borrow to rebuild, credit demand rises, and whether that demand is met depends on the condition of the local lender. Access to credit tightens after a disaster, but borrowers with an established rela-

tionship are largely shielded (Berg and Schrader, 2012), and areas served by local lenders see credit expand and recover faster (Cortés, 2014). Banks whose own balance sheets are damaged transmit that damage into the investment of the firms they finance (Hosono et al., 2016), well-capitalized independent banks rebuild capital and reallocate their portfolios after a catastrophe (Schüwer et al., 2019), and regional lenders exposed to a flood expand recovery lending to affected borrowers (Koetter et al., 2020). Public guarantees change this allocation by letting banks extend credit into the affected area while shifting the risk off their own books (Becerra-Ornelas et al., 2026). All of this describes the disaster zone. It says little about the rest of the lender's portfolio.

Propagation beyond the disaster zone is where the systemic question lies, and the evidence there is thinner. The clearest results are for mortgage markets. Lenders active in a disaster region reallocate funds toward it to meet the surge in local demand and cut mortgage credit in unaffected areas (Cortés and Strahan, 2017), and the same cross-region spillover through branch networks appears after El Niño shocks (De Marco and Limodio, 2024). Where the shock is followed through an entire lending network, the event is gradual rather than extreme, as with persistent drought in Brazil, where lending first flows into the affected region and later contracts network-wide once borrowers default (Mestieri et al., 2026). For corporate loan markets after an extreme event, the evidence is either about prices, showing that spreads rise for at-risk but undamaged borrowers in a way that tracks the salience of the storm (Correa et al., 2026), or about quantities identified at the county level, comparing lenders that lend into the same county (Ivanov et al., 2022). That design controls for demand shocks common to a county, and it is the closest paper to ours. Our evidence is one level more granular. We compare lenders within the same firm, so fixed effects absorb the borrower's own credit demand rather than its county's, and we test a heterogeneity their design does not, the hurricane risk of the borrower's own location.

This paper asks how do local physical climate shocks propagate through financial intermediaries in corporate loan markets? We answer it using Hurricane Harvey, which struck Texas and Louisiana in August 2017, and U.S. syndicated loan data from DealScan. Our treatment is the lender's exposure to the disaster, measured as the share of its pre-storm syndicated portfolio lent to borrowers in counties that subsequently received FEMA major disaster declarations. Building this measure requires knowing how much of each syndicated loan each lender holds, and DealScan reports individual lender shares for only part of the

market. We therefore follow Chodorow-Reich (2014), keeping every reported share and imputing the missing ones from the average shares reported by lenders holding the same role in syndicates of the same structure.

In August 2017, Hurricane Harvey caused widespread damage to counties in Texas and Louisiana, increasing the demand for credit. We first document this direct effect in our own data. Using county-level damages,[1] borrowers headquartered in damaged counties increased their total syndicated borrowing by about 12 percent relative to other borrowers between the year before the storm and the year after, and the increase rises with the dollar value of damage in the borrower's county. This is a correlation rather than a causal estimate, but it is a stable one. Controls for the borrower's pre-storm credit, its number of lenders, its industry, and its state leave the pattern intact, and it matches the direct effect that Ivanov et al. (2022) document across disasters. The disaster zone did not lose credit after Harvey. It gained credit. Moreover, lenders with different Harvey exposure expanded credit to the same affected borrower at the same rate, and exposed lenders held their market share in the damaged counties. The reallocation this paper documents happened outside the disaster zone.

To separate supply from demand, the central difficulty in any study of credit after a disaster, we use the within-borrower estimator of Khwaja and Mian (2008). Because most syndicated borrowers hold credit from several lenders at once, 92 percent of tranches in our sample have more than one lender parent and the median tranche has four, we can compare, within the same unaffected borrower, the change in credit from lenders with different Harvey exposure between the year before the storm and the year after. Borrower fixed effects absorb any shock to the borrower's demand for credit, so the remaining variation across its lenders reflects supply. The design requires only that credit demand be common across a borrower's lenders.

Our first finding is that a geographically concentrated disaster transmits to borrowers the storm never touched as a contraction of credit supply, and that the contraction operates through quantities. A borrower's credit from a lender one standard deviation more exposed to Harvey, 6.5 percentage points of portfolio share, falls by 4.8 percent relative to its other

[1] County-level damage data come from the Spatial Hazard Events and Losses Database for the United States (SHELDUS), Arizona State University. We measure Harvey damages as flooding and hurricane damages recorded in August and September 2017 in the counties that received FEMA major disaster declarations for Harvey.

lenders, weighting each lending relationship by its pre-storm loan amount so that the estimate describes dollars of credit, and by 4.0 percent when borrowers hit by Hurricane Irma are also excluded. Counting each relationship equally instead, the two estimates are 3.4 and 2.4 percent, so the cuts fell disproportionately on the larger loans. Spreads and maturities show no adjustment. The result is not driven by borrowers near the storm. Excluding all borrowers within 300 or 700 miles of the Harvey impact zone leaves the estimate essentially unchanged, at 4.7 and 4.3 percent. Whatever links these distant borrowers to Harvey operates through the lender.

Our second finding is that the contraction targets climate risk rather than geography. Exposed lenders do not withdraw from a region. They withdraw from a type of borrower. Splitting unaffected borrowers by their county's historical hurricane frequency, the supply cut loads on the risky side of the split. Borrowers in counties with at least ten historical hurricane passages lose 7.4 percent of loan volume per standard deviation of exposure, while borrowers in low-risk counties are statistically unaffected. The same reallocation is visible at the level of the market. A one standard deviation more exposed lender moved about two percentage points of its portfolio out of hurricane-risky counties, off a pre-storm base of 65.8 percent, and ceded market share in risky counties while holding it in safe ones. What looks from inside a lending relationship like a credit cut is, at the level of the market, a transfer of climate risk across lenders.

Our third finding concerns the sequence of adjustment along the intensive and extensive margins. Amounts are cut first, in the two to three quarters after the storm. Then, two to four quarters later, exposed lenders begin to let lending relationships lapse. The probability that a pre-existing relationship remains active falls by 1.1 percentage points per standard deviation of exposure by the end of 2019 when relationships are counted equally, and the lapse is concentrated among high-hurricane-risk borrowers. The intensive-margin gradient in amounts reappears in who loses the relationship. The relationships that lapse are disproportionately small ones, so the extensive margin is visible in relationship counts more than in dollars.

The response is specific to Harvey. Re-running the entire design around placebo event dates in the third quarters of 2014 and 2015 produces coefficients that are uniformly small and insignificant. A third placebo date, the third quarter of 2016, is not a clean falsification test, because its post window runs through September 2017 and straddles the recovery of

the oil market. There, exposed lenders expand rather than contract, the opposite sign of a supply cut. Pre-storm concentration in the Harvey counties predicts lending cuts around Harvey and lending growth when oil recovers, so the contraction is a response to the event, not a stable trait of Texas-concentrated lenders.

Why would a lender respond to a disaster by cutting credit to borrowers the disaster did not touch? The timing offers an interpretation. Harvey was one of the costliest disasters in U.S. history and arrived precisely as climate risk entered the financial stability agenda. The Financial Stability Board's Task Force on Climate-related Financial Disclosures published its recommendations in June 2017, and the Network for Greening the Financial System was founded in December 2017 (Network for Greening the Financial System, 2018). Google search interest in climate change around Harvey was roughly two and a half times its level around Hurricane Sandy five years earlier. Salient events shift managers' risk assessments even when fundamentals are unchanged (Dessaint and Matray, 2017), and salience-driven repricing after disasters has been documented in loan spreads (Correa et al., 2026). Our evidence is consistent with the quantity counterpart. A salient shock made latent hurricane risk visible, and the lenders holding the most of it rebalanced away from hurricane-risky borrowers throughout their portfolios. Climate risk was not eliminated. It was reallocated.

To our knowledge, these are the first estimates of a credit supply reallocation in corporate loan markets after an extreme weather event that control for credit demand at the level of the individual borrower. The within-borrower design compares lenders with different disaster exposure lending to the same unaffected firm, so the estimates separate supply from demand, the central difficulty in the literature that follows disasters into credit markets.

Related Literature. This paper contributes to three strands of the literature. First, it relates to work documenting that physical climate risk is priced in financial markets and credit contracts (Bernstein et al., 2019; Ouazad and Kahn, 2022; Goldsmith-Pinkham et al., 2023; Javadi and Masum, 2021). Javadi and Masum (2021) show that firms in locations with higher long-run climate-change exposure pay significantly higher loan spreads, with effects concentrated among long-maturity loans to poorly rated firms, suggesting lenders increasingly price physical climate risk into contract terms. Most closely related, Correa et al. (2026) show that loan spreads increase for at-risk but not directly affected borrowers following climate-linked disasters. We study the same borrower margin, but not with the

same source of variation. They hold the lender fixed with bank-by-year effects and compare two borrowers of one bank, which recovers how the market reprices a borrower once a disaster makes its risk salient. We hold the borrower fixed and compare two lenders of one firm, which recovers how much of the response is specific to the lenders the storm reached. Importantly, they attribute the repricing primarily to *salience bias*, since increased spreads are short-lived, amplified by media attention, and driven by lenders overreacting to the prominence of the event rather than by a permanent reassessment of long-run credit risk (Dessaint and Matray, 2017). Our contribution is the quantity counterpart, identified causally. Exposed lenders cut credit to unaffected borrowers relative to the same borrower's other lenders, the cut sorts on the borrower's own hurricane risk rather than on proximity to the storm, and the adjustment operates through amounts and relationship terminations rather than through spreads.

Second, we contribute to the literature on bank geographic diversification and spatial credit spillovers (Goetz et al., 2016; Chu et al., 2020; Conlon et al., 2024; Ivanov et al., 2022; De Marco and Limodio, 2024; Mestieri et al., 2026). A canonical mechanism emphasizes internal capital markets, whereby banks reallocate lending across regions in response to shocks (Cortés and Strahan, 2017; Koetter et al., 2020). Ivanov et al. (2022) show, using syndicated loan data, that when disaster areas experience increased credit demand, banks in their networks cut lending to borrowers in distant unaffected regions by approximately 3 percent, a pattern directly analogous to the credit supply contraction we document. De Marco and Limodio (2024) document a similar cross-county spillover for mortgage lending following El Niño events, establishing that even non-destructive climate shocks propagate through bank branch networks to counties not directly affected by temperature increases. Rehbein and Ongena (2022) show that low bank capital amplifies local shock spillovers. Firms in non-disaster regions connected to disaster-exposed banks reduce borrowing and employment, with retrenchment largest in areas the bank anticipates as future disaster risks. At the macro level, Mestieri et al. (2026) study how persistent drought shocks in Brazil propagate through bank branch networks. In the short run banks channel capital *into* drought-affected regions, but when shocks persist and borrowers default, lending contracts throughout the network, with spillovers to indirectly connected regions roughly half the size of the direct effect. We study a complementary channel, how lenders' geographically distributed portfolio exposures translate into differential credit allocation in integrated lending markets, including for borrowers in unaffected locations.

Third, the paper builds on the quasi-experimental literature that identifies credit supply by comparing lenders within the same borrower. Khwaja and Mian (2008) introduce the within-borrower estimator using liquidity shocks to Pakistani banks. Jiménez et al. (2012) apply the logic to loan applications in Spain, Peek and Rosengren (2000) trace shocks to Japanese banks into U.S. real activity through their local lending, and Ivashina and Scharfstein (2010) document the contraction of syndicated lending during the 2008 crisis. Closest in method, Chodorow-Reich (2014) uses lenders' differential exposure to the Lehman collapse to identify credit supply in the syndicated market and trace it into employment, and provides the syndicate-share imputation we adopt. We bring this identification toolkit to an extreme weather event. Lender exposure to Hurricane Harvey plays the role of the funding shock, and the within-borrower comparison isolates how the reassessment of climate risk is transmitted to borrowers the storm never touched.

The remainder of the paper proceeds as follows. Section 2 describes the data, covering the hurricane exposure of U.S. counties, the syndicated loan market, the construction of the sample, and its composition. Section 3 sets out a portfolio model of lender reallocation that separates a loss of lending capacity from a revision in beliefs about hurricane risk, and identifies the coefficient that distinguishes them. Section 4 presents the empirical strategy, describing how we measure lender exposure to Harvey and how we estimate credit supply, with the identifying assumptions. Section 5 first documents the credit expansion inside the disaster zone and how it was financed, and then reports the results for the intensive and extensive margins, the robustness checks, and the interpretation. Section 6 concludes.

2 Data

Hurricane Harvey made landfall on the Texas coast on August 25, 2017 as a Category 4 storm, then stalled over the Houston metropolitan area for several days and produced record rainfall and catastrophic flooding across Texas and Louisiana. With estimated damages of about $125 billion it is the second costliest natural disaster in United States history, behind only Hurricane Katrina.[2] We use Harvey as a case study of an extreme climate event. This section describes the data. We first build each county's hurricane risk from the historical

[2]Damage figures are from the National Oceanic and Atmospheric Administration's billion-dollar weather and climate disasters record (NOAA, 2025).

storm record and define the area Harvey affected, using FEMA disaster declarations for its footprint and SHELDUS damages for its intensity. We then describe the syndicated loan market, the imputation of lending shares within each syndicate, and the construction of the estimation sample.

2.1 Hurricane frequency and damages

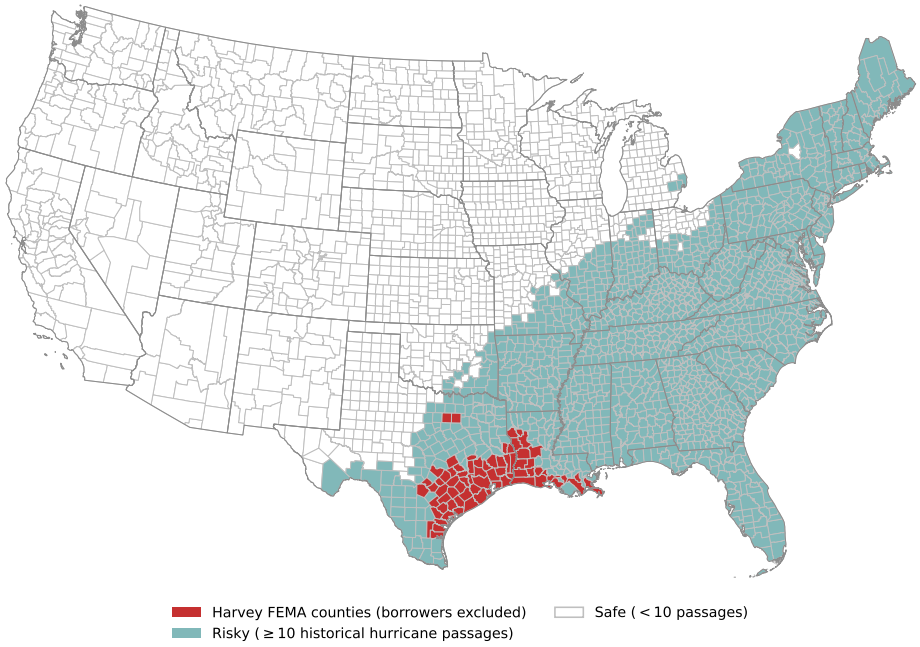
The mechanism we test is that exposed lenders rebalance away from hurricane-prone borrowers everywhere, not just near Texas. This requires a measure of each borrower's own hurricane risk. For every U.S. county we count the number of hurricanes that passed within 100 miles of the county centroid between 1851 and the day before Harvey's landfall, using NOAA's HURDAT2 storm-track database.[3]

Table A1 in Appendix A characterizes the distribution of this measure across the 3,233 counties of the continental United States. Hurricane exposure is widespread but very unequal. The median county saw 13 hurricane passages within 100 miles over the 166-year record, one quarter of counties saw at most one, and the most exposed counties, in southern Florida, saw more than 100. The same skewness appears in close passages (within 50 miles). Panel B reports the measure across the borrowers in our sample. Because economic activity, and with it syndicated borrowing, concentrates along the coasts, sample borrowers sit in more exposed counties than the average county. The median borrower's county saw 18 hurricane passages, 65 percent of borrowers are in counties with at least ten, the threshold we later use to define risky counties, and only 12 percent are in counties with none.

Figure 1 maps the design. It shows the Harvey counties dropped from the estimation sample and the split of the remaining counties into risky and safe at ten historical hurricane passages, the threshold used throughout the paper. Hurricane risk is concentrated along the Gulf and Atlantic coasts.

[3]HURDAT2 is the National Hurricane Center's revised Atlantic hurricane database, the official post-season reanalysis of every Atlantic tropical cyclone since 1851. For each storm it records position, maximum sustained winds, and central pressure at six-hourly intervals. We count a passage when the storm is at hurricane intensity, maximum sustained winds of at least 64 knots, at the time it comes within 100 miles of the county centroid. The database is publicly available at https://www.nhc.noaa.gov/data/hurdat/.

Figure 1: Design. Dropped Counties and Risky Versus Safe Counties



Notes: Harvey FEMA-declared counties (dropped from the estimation sample) are shown in red. The remaining counties are split into risky (at least 10 historical hurricane passages within 100 miles of the county centroid since 1851, NOAA HURDAT2) and safe (<10 passages). This is the split used in the risk interaction (Table 6), the market-share analysis (Table 7), the event-study split (Figure A2), and the exit split (Figure 5).

To identify the timing and geographic scope of the shock we use FEMA major disaster declarations, which determine eligibility for federal assistance and provide an administrative, storm-specific definition of the affected area.[4] Combining the historical HURDAT2 exposure with the FEMA declarations separates long-run vulnerability, which every coastal borrower carries, from the realized 2017 shock, which only Harvey counties experienced. Where the analysis requires the intensity of the shock rather than its footprint, we measure the dollar value of Harvey damage in each affected county with SHELDUS.[5]

[4]FEMA disaster declarations DR-4332 (Texas, declared August 25, 2017) and DR-4345 (Louisiana). Under the Stafford Act, a major disaster declaration is issued by the President at the request of the state's governor and activates federal individual and public assistance programs in the designated counties. The county designations are published in FEMA's disaster declarations database at https://www.fema.gov/disaster/declarations.

[5]The Spatial Hazard Events and Losses Database for the United States (SHELDUS), maintained by the Center for Emergency Management and Homeland Security at Arizona State University, records county-

2.2 Syndicated loans

Loan-level data come from Refinitiv DealScan, which covers U.S. syndicated credit and reports tranche amounts, spreads over the benchmark rate, maturities, the identity and location of the borrower, and the composition of the lending syndicate.[6] In a syndicated loan, a group of lenders jointly extends credit to a single borrower under one agreement, coordinated by one or more lead arrangers who structure the deal and negotiate its terms, while participant lenders fund shares of it. The unit of observation is the tranche, a credit facility within a loan agreement. One agreement may contain several tranches that differ in type, pricing, and maturity but are issued simultaneously to the same borrower. Because tranches are funded jointly, most borrowers in this market have several simultaneous lenders, which is precisely the variation the within-borrower design of Section 4 exploits.

Syndicated loans are floating-rate instruments. The borrower pays a short-term benchmark rate plus a fixed credit spread, the margin, which is set at origination and is the object of interest in our spread regressions. The benchmark is the London Interbank Offered Rate (LIBOR) for most of our window and Term SOFR, the Secured Overnight Financing Rate published by the CME Group, for tranches originated after the LIBOR transition began. Because the benchmark moves with monetary policy and is common to all borrowers, the margin is the part of the price a lender sets, which is why we measure repricing by the change in the margin rather than by the change in the all-in rate.

We keep tranches whose borrower is located in the United States with a valid five-digit zip code, whose deal status is neither cancelled nor rumored, and which are denominated in U.S. dollars. We further require a non-missing tranche amount, margin, and maturity date, and we restrict the sample to tranches priced off LIBOR or Term SOFR so that spreads are measured on a comparable basis. Lenders are aggregated to the parent institution, so that divisions and affiliates of the same banking group are treated as one lender.[7] Each borrower

level dollar losses, injuries, and fatalities by hazard type and month since 1960, compiled primarily from the National Centers for Environmental Information storm records. We measure Harvey damage as the flooding and hurricane damage recorded in the FEMA-declared counties in August and September 2017. The database is available at `https://cemhs.asu.edu/sheldus`.

[6] DealScan "collects loan-level information on syndicated loans from Securities and Exchange Commission (SEC) filings, company statements, and media reports, and attempts to process the universe of such loans" (Chodorow-Reich, 2014).

[7] A single organization participates in DealScan under several legal entities. JPMorgan Chase & Co., for example, appears as JPMorgan Chase Bank N.A., J.P. Morgan Securities, and Chase Bank USA, among

is located at its headquarters, whose zip code is mapped to a county using the Department of Housing and Urban Development (HUD) zip-to-county crosswalk.[8]

Lending shares. DealScan reports each lender's share of a tranche for only a small fraction of syndicates, under 4 percent of tranches in our window. To attribute dollar amounts to individual lenders we follow Chodorow-Reich (2014). In fully reported deals, the share a lender takes depends systematically on its role and on the structure of the syndicate, since lead arrangers retain larger shares than participants and both shares shrink as the syndicate grows. We therefore use the fully reported tranches to compute the average share of a lead arranger and of a participant within every syndicate-structure cell, defined by the number of leads and the number of participants, and use these averages to fill in missing shares. Among fully reported deals with one lead arranger and one participant, for instance, the lead retains 62 percent of the tranche on average and the participant 38 percent, so those are the weights applied to deals of that structure whose shares are missing. Reported shares are never discarded. When all shares are reported we renormalize them to sum to one. When none are reported, each lender receives the average share for its role and syndicate structure. In mixed cases, reporting lenders keep their reported shares and the unreported residual is allocated across the remaining lenders in proportion to the role-by-structure averages. A lender's amount on a tranche is its share times the tranche amount.

The main results do not depend on this imputation rule. Section 5.5 re-estimates the pooled specification keeping the reported shares and imputing the missing ones by an equal split among the non-reporting lenders, and by each non-reporting lender's size in the pre-storm syndicated loan market (Table A11 in Appendix E). The estimated contraction is 3.0 and 2.5 percent per standard deviation of exposure under the two rules, against 3.4 percent in the baseline.

The sample window is one year before Harvey against one year after. The pre-period runs from July 2016 through June 2017 and the post-period from October 2017 through September 2018. The hurricane quarter itself, the third quarter of 2017, is dropped, so

others. We roll subsidiaries up to their ultimate parent using the DealScan parent identifier before computing portfolio shares, so a lender's exposure reflects the consolidated organization's U.S. syndicated portfolio.

[8]Locating DealScan borrowers at their headquarters is standard in the climate-finance literature. Javadi and Masum (2021) measure a firm's climate risk exposure by the state of its headquarters, citing evidence that headquarters sit close to firms' operations and core business activities. The proxy is weakest for firms with geographically dispersed operations, which we return to in Section 5.5.

the comparison is between complete pre- and post-storm years that exclude the disruption window. The resulting sample contains 26,619 lender-borrower relationships involving 8,800 tranches, 4,118 borrowers, 1,144 lenders, and 662 lender parents. At the tranche-lender level the sample has 51,554 observations.

Table A2 in Appendix A shows how the 26,619 relationships split across margins. About one fifth lend in both years, half lend only in the pre-year, and a third appear only in the post-year. This turnover reflects observing syndicated credit at origination. A relationship is recorded only when a new tranche is signed, so a pair with no originations in a window is absent from that window even if loans remain outstanding. The 5,512 pairs that lend in both years form the intensive margin, where the within-borrower comparison is available, and they are the estimation sample for the cross-sectional results. The extensive margin, the exit and entry of relationships, is analyzed in Section 5.3. The margin shares are essentially identical when Harvey borrowers, or Harvey and Irma borrowers, are dropped.

Selection into the intensive margin could itself respond to the shock. It does not work against us. In a linear probability model of ending the relationship, pre-only against both margins, on standardized Harvey exposure, the coefficient is -0.023 with a standard error of 0.009, so more exposed lenders are slightly *less* likely to drop out of the intensive-margin sample, and the sample is not selectively losing exposed lenders' relationships. If anything this selection is conservative for the contraction result. Entry, post-only against both, is unrelated to exposure.

Table A3 in Appendix A reports the composition of the two sides of the market, using DealScan's institution classifications. Banks dominate the lender side. U.S. banks are 41 percent of the 1,144 lenders but 58 percent of tranche-lender observations, and together with foreign banks they account for 69 percent of lenders and 86 percent of observations. The remaining third of the lender population, made up of finance companies, investment banks, direct lenders, investment funds, and insurers, holds only 14 percent of observations. Nonbanks are numerous but each participates in far fewer syndicates. This asymmetry matters for the results, because the reallocation we document is concentrated among the large bank lenders that account for most of the market. On the borrower side, corporates account for 82 percent of the 4,118 borrowers, followed by nonbank financial firms at 9 percent and utilities at 5 percent. At the detailed industry level the sample is diversified. The largest group, financial services, holds 9 percent of borrowers, manufacturing, healthcare, and

technology hold 6 to 8 percent each, and oil and gas holds 6 percent. No industry dominates, and in particular the energy sector's direct link to the Gulf Coast cannot drive the results.

Table A4 in Appendix A reports sample composition by loan type and purpose, at the level of unique tranches and of tranche-lender observations, together with the structure of the underlying records. Term loans account for 52 percent of unique tranches and revolvers for 43 percent. At the observation level the ranking flips, 45 against 51 percent, reflecting the larger syndicates behind revolving facilities. By purpose, general-purpose loans constitute two thirds of tranches and nearly three quarters of observations, followed by acquisition financing, leveraged buyouts, and takeovers. Acquisition-related purposes together account for about a fifth of tranches. Panel C describes the records themselves. The panel tracks lending activity at origination and at amendment, and 73 percent of tranche-lender activity records are amendments of existing tranches rather than new issuances, so restricting attention to originations alone would miss most of the activity in the market.

Table A5 in Appendix A characterizes the whole sample, intensive and extensive margins together, separately for the pre-year and the post-year. Three features stand out. First, the composition of the market is stable across the storm. U.S. banks account for about 55 percent of lender appearances in both windows and the term-loan and revolver mix is unchanged. Second, aggregate credit conditions did not tighten after Harvey. The median margin falls from 300 to 275 basis points, the median tranche grows from \$140 to \$150 million, and the median tenor shortens only slightly. Whatever adjustment the storm triggered is not visible as an aggregate credit crunch. Third, the share of lending flowing to borrowers in Harvey counties is essentially unchanged at 12.5 percent of volume in both windows, a first indication that the aggregate map of credit conceals the reallocation underneath. Syndication is the norm throughout, with 92 percent of tranches having more than one lender parent and the median tranche having four, which is what makes the within-borrower design feasible.

The stability of aggregate conditions is not specific to borrowers outside hurricane-prone areas. Table A6 in Appendix A repeats the pre–post comparison splitting borrowers by their county's historical hurricane risk, using the same threshold as the heterogeneity results below, at least ten hurricane passages within 100 miles of the county since 1851. Conditions move in parallel in the two groups. The median margin falls from 300 to 275 basis points for hurricane-prone and less-prone borrowers alike, median tranche amounts rise in both groups, and median tenors shorten by one to three months. Hurricane-prone borrowers are

not being repriced or rationed in the aggregate after Harvey, which anticipates the central finding. The adjustment to the storm is a reallocation across lenders within borrowers, invisible in market-wide terms.

3 A Portfolio Model of Lender Reallocation

This section writes down a two-period portfolio problem to interpret the empirical results and mechanisms. Consider a lender that cuts credit after a disaster, this can be explain by being short of capacity, having absorbed losses in the damaged region, in which case it lends less to everyone. Or it can be explained by having revised its assessment of hurricane risk, in which case it lends less to the borrowers it now considers exposed. In a pooled estimate the two are observationally identical, since both produce a negative coefficient on lender exposure.

The model is the mean-variance portfolio problem of Markowitz (1952) applied to a financial intermediary, following Pyle (1971) and Hart and Jaffee (1974), and closed with a capacity constraint in the manner of Holmström and Tirole (1997), in which the intermediary's own capital limits how much it can lend. Two elements of that literature are what we need. First, an intermediary facing costly external finance behaves as if risk averse, so it prices a marginal loan by that loan's covariance with the portfolio it already holds rather than by the loan's own variance (Froot et al., 1993; Froot and Stein, 1998). Second, when the risk being borne is catastrophe risk, the capital available to bear it is limited and its shadow price moves with recent losses (Froot, 2001). We add to this a belief about the hurricane factor that is estimated rather than known, which is the ingredient that makes portfolios respond to news about a risk or changes in the financial regulation, and not only to news about capital.

Environment. There are two periods and a continuum of borrowers indexed by i. Borrower i sits in a county with hurricane risk h_i, the historical frequency of hurricane passages defined in Section 2. A loan to i returns

$$R_i = r_i - \lambda h_i x + \varepsilon_i, \tag{1}$$

where r_i is the contracted spread, x is an aggregate hurricane realization with mean θ and variance σ^2, λ scales the loss per unit of risk, and ε_i is idiosyncratic with variance s^2. The term $\lambda h_i x$ is what makes hurricane-prone borrowers different. Their returns load on a common climate factor, so they are correlated with each other in a way that low-h borrowers are not.

The lender's problem. Lender j solves the same portfolio problem in each of the two periods, taking as given its current beliefs θ_j about the hurricane factor and its current capital K_j. It picks the portfolio $\{L_{ij}\}$ that maximizes the expected return on its loan book minus a penalty on the variance of that book, subject to a capacity constraint that caps total lending at its available capital K_j:

$$\max_{\{L_{ij}\}} \sum_i L_{ij}\, E_j[R_i] - \frac{\gamma}{2} \operatorname{Var}_j\Big(\sum_i L_{ij} R_i\Big) \qquad \text{subject to} \qquad \sum_i L_{ij} \le K_j. \tag{2}$$

The problem is static and repeated. Loans last one period, so the lender rebuilds its book from scratch each period. In period 0 every lender holds the common prior θ^0 and its pre-storm capital. Harvey then changes both inputs, and the period-1 problem is the same maximization with the updated beliefs and capital. Nothing else links the periods, so the model's predictions concern the difference between the two portfolios, which is the object the pre–post regressions estimate. The parameter $\gamma > 0$ is the lender's effective risk aversion, the penalty it puts on portfolio variance. The constraint limits total dollars lent, not risk-weighted dollars, so a dollar to a Houston borrower and a dollar to a Minneapolis borrower use the same capacity. Letting $\mu_j \ge 0$ be its multiplier and $H_j \equiv \sum_k L_{kj} h_k$ the lender's aggregate climate exposure, the first-order condition, derived step by step in Appendix B, gives

$$L_{ij} = \frac{r_i - \mu_j - \lambda \theta_j h_i - \gamma \lambda^2 \sigma^2 H_j h_i}{\gamma s^2}. \tag{3}$$

The term $\gamma \lambda^2 \sigma^2 H_j h_i$ is the covariance charge of Froot and Stein (1998), under which a loan to a hurricane-prone borrower is expensive to a lender already holding hurricane risk, and cheap to one that is not.[9] The shadow cost of capacity μ_j enters alone, so a tighter constraint

[9]The covariance pricing of Froot and Stein (1998) applies only to risks the bank cannot hedge in a liquid market, and hurricane risk inside a corporate loan portfolio is that case, since no traded instrument pays off on it. In their model risk aversion also rises as capital falls, so strictly Harvey should raise exposed lenders'

shifts lending down by the same amount for every borrower. Beliefs θ_j enter only multiplied by h_i, so a revision in beliefs tilts the portfolio against hurricane-prone borrowers and leaves everyone else untouched.

Harvey. The storm consumes capacity, lowering K_j by more for lenders with more of their portfolio in the damaged counties, and it is news about the hurricane factor. Lenders update

$$\theta_j^1 = \theta^0 + \kappa(e_j)\,\Delta, \qquad \kappa'(e_j) > 0, \tag{4}$$

where e_j is the lender's pre-storm Harvey exposure of equation (7), $\Delta > 0$ is the surprise, and κ governs how much weight the lender puts on it. Exposed lenders update more because the storm hit their own book, an availability effect documented for managers by Dessaint and Matray (2017). A Bayesian observing the same public storm would put the same weight on it regardless of exposure, so $\kappa'(e_j) > 0$ is a departure from the common-prior benchmark rather than an implication of it.

In equation (4), Δ is the news that moves the lender's assessment of hurricane risk. It can come from the storm itself or from the supervisory push for climate disclosure, since the TCFD framework appeared two months before landfall. We cannot separate the two, because both operate through the same margin. A lender with more Harvey-county loans learned more from the storm and also had more climate exposure to disclose. Nor do we need to. Both are revisions of θ_j, and δ_H is built to distinguish a revision of θ_j from a fall in K_j, not to say which force prompted the revision.

What each channel predicts. Substituting (4) into (3) and differencing between the pre- and post-storm year, step by step in Appendix B, gives

$$\Delta \ln L_{ij} \;\approx\; \alpha_i - \phi(e_j) - \psi(e_j)\,h_i, \qquad \phi' > 0, \quad \psi' > 0, \tag{5}$$

γ as well as lower K_j. We hold γ fixed because equation (3) shows a higher γ only scales down lending and damps the tilt on h_i, so it cannot manufacture the negative interaction we test for, and if anything it works against finding one.

where ϕ collects the capacity term and ψ the belief term.[10] Equation (5) is the estimating equation (13) with α_i the borrower fixed effect, so

$$\beta = -\left[\phi' + \psi' \bar{h}_L\right], \qquad \delta_H = -\psi' (\bar{h}_H - \bar{h}_L), \tag{6}$$

with $\bar{h}_H$ and $\bar{h}_L$ the average risk in the high- and low-risk groups, and the algebra behind the mapping in Appendix B. A pure capacity shock sets $\psi' = 0$, so the contraction is uniform, with $\beta < 0$ and $\delta_H = 0$. Belief updating delivers $\delta_H < 0$, and it also implies that β is small and negative rather than exactly zero, because even low-risk borrowers carry some h_i. The heterogeneity coefficient, not the pooled one, is what distinguishes the channels. The identification has a mirror image in Froot (2001). After Hurricane Andrew, reinsurance prices rose for earthquake risk as well, a risk the storm carried no news about, which identifies scarce capital. Our test reads the tilt on h_i, the dimension a pure capital shock is silent on, as the signature of beliefs.

4 Empirical Strategy

This section presents the empirical strategy that recovers the two coefficients of equation (6). We first construct the continuous lender treatment, the share of each lender's pre-storm syndicated portfolio lent in the counties Harvey damaged. We then present the within-borrower estimator that separates credit supply from credit demand, together with its identifying assumption and the choice of dollar weights. The results section applies this toolkit first inside the disaster zone, where borrowing rose, and then to the rest of the United States, on the intensive margin of loan amounts, on its heterogeneity in the borrower's own hurricane risk, and on the extensive margin of relationship exits. In parallel to these pair-level designs, a lender-level design in portfolio and market shares asks whether the reallocation is large enough to move the aggregate geography of credit.

[10] Two approximations. First, the constraint binds on dollars while the outcome is a log change, so a uniform dollar cut is a larger proportional cut at smaller relationships. Table A15 shows that high-risk pairs are the larger ones, so this biases δ_H toward zero. Second, we hold the feedback through H_j fixed. In equilibrium a lender that sheds hurricane-prone loans also lowers its own H_j, which lowers the covariance charge and damps the tilt, so the feedback biases δ_H toward zero rather than away from it.

4.1 Lender exposure to Harvey

The treatment variable is the lender's pre-storm portfolio concentration in the counties Harvey hit,

$$\text{Harvey exposure}_j = \frac{\text{lender } j\text{'s syndicated loan volume in Harvey counties}}{\text{lender } j\text{'s total U.S. syndicated loan volume}}, \tag{7}$$

measured over the pre-period at the lender parent level, where Harvey counties are the Texas and Louisiana counties under FEMA major disaster declarations. In the estimation sample the pair-weighted mean exposure is 13.4 percent with a standard deviation of 6.5 percentage points, and the parent-level mean is 10.9 percent. Most lenders hold under 20 percent of their portfolio in Harvey counties, with a long right tail of locally concentrated lenders (Figure A1 in Appendix C). Throughout the paper exposure is standardized to unit variance, so coefficients measure the effect of a one standard deviation (6.5 percentage point) increase in exposure.

Borrowers located in Harvey FEMA counties receive 10 percent of tranches and 12 percent of dollar volume in the pre-period, shares that are stable before the storm. All results in this paper concern the other 90 percent of the portfolio, the borrowers in counties Harvey never touched.

Splitting lender parents at the median of Harvey exposure and comparing their pre-period lending shows that exposure is not a small-lender trait. Above-median-exposure lenders are the large, central players in the market. Their pre-period portfolios are two orders of magnitude larger ($10.3 billion against $102 million), they lend to 64 borrowers against 2, act as lead arranger more often (36 against 22 percent), tilt toward revolvers, and price cheaper (281 against 331 basis points). The supply contraction documented below therefore comes from large, mainstream, low-margin lenders rather than from marginal or high-priced ones.[11]

[11] The extreme upper tail of the exposure distribution does consist of small local lenders. The ten most exposed parents include single-loan Texas banks with exposure near one. The market-share estimates of Section 5.4, which weight lenders by pre-period volume, show that the aggregate reallocation is carried by large lenders.

4.2 Borrowing inside the disaster zone

Harvey destroyed property on a large scale, and rebuilding requires financing, so credit demand in the damaged counties should have risen. Before we can read anything as a supply response elsewhere, we need to know what happened inside the disaster zone. We measure it with the cross-sectional regression

$$\Delta \ln(\text{credit})_i = \alpha + \beta \cdot \text{Damage}_{c(i)} + \gamma' X_i + \varepsilon_i, \tag{8}$$

where $\Delta \ln(\text{credit})_i$ is the change in borrower i's total syndicated credit between the pre-period and the post-period, and Damage is measured in the borrower's headquarters county, either as an indicator for positive Harvey damage or as the log of one plus the dollar value of damage. This specification follows the direct-effect regression that Ivanov et al. (2022) estimate across natural disasters. It compares damaged-county borrowers to all other borrowers, so β is a correlation rather than a causal estimate. Borrowers in and out of the disaster zone can differ for many reasons. The point of Table 1 is that the correlation is stable as each column adds controls, from a regression with damage as the only regressor to one that also holds the borrower's pre-storm credit, its number of lenders, its industry, and its state.

4.3 Credit supply

The obstacle to measuring credit supply is that lending is an equilibrium quantity. When lending to a borrower falls, the borrower may have stopped asking or the lender may have stopped offering, and the two are observationally identical in loan data. Khwaja and Mian (2008) solve this by using borrowers that hold more than one lending relationship. Let the change in credit from lender j to borrower i be

$$\Delta \ln(\text{amount})_{ij} = \underbrace{\alpha_i}_{\text{demand}} + \underbrace{\eta_j}_{\text{supply}} + \varepsilon_{ij}, \tag{9}$$

where α_i collects everything moving the borrower's demand for credit and its creditworthiness, and η_j collects everything moving the lender's willingness to extend it. Both terms are unobserved, and in a single-lender borrower they cannot be told apart. In a borrower with two lenders they can. Differencing equation (9) across two lenders of the same borrower

removes α_i entirely, whatever it contains, and leaves $\eta_j - \eta_{j'}$. Demand does not have to be modeled or measured. It has to be common across the borrower's lenders, and then it cancels.

We take Harvey exposure as the observable component of η_j and estimate

$$\Delta \ln(\text{amount})_{ij} = \alpha_i + \beta \cdot \text{Harvey exposure}_j + \varepsilon_{ij}, \tag{10}$$

where $\Delta \ln(\text{amount})_{ij}$ is the change in lender j's credit to borrower i between the post-year and the pre-year, α_i is a borrower fixed effect, and Harvey exposure is the standardized portfolio share from equation (7). Standard errors are clustered at the lender parent. We also estimate the same specification for the change in the loan spread and in maturity. Identification comes from comparing two lenders with different Harvey exposure lending to the same unaffected borrower, so β is the credit response to a one standard deviation increase in exposure, holding the borrower fixed.

Equation (10) compares one year to one year and delivers a single number. To see when the adjustment happens, and to check that exposed and unexposed lenders were on the same path before the storm, we estimate a dynamic version on a quarterly lender-borrower panel,

$$\ln(\text{amount})_{ijt} = \alpha_{it} + \gamma_j + \sum_{k \neq -1} \beta_k \, \mathbf{1}[t = k] \cdot \text{Harvey exposure}_j + \varepsilon_{ijt}, \tag{11}$$

where α_{it} are borrower-quarter fixed effects, γ_j are lender fixed effects, and the omitted period is the quarter before landfall (2017Q2). The borrower-quarter fixed effect does the same work as α_i above, quarter by quarter, so each β_k is the lending gap between a more and a less exposed lender serving the same borrower in the same quarter. The coefficients for $k < -1$ are the test of common trends, and those for $k \geq 0$ trace the timing and persistence of the response.

Two features of this estimator shape how the results should be read. First, β measures a lender's lending relative to the other lenders of the same borrower, not the total credit that borrower receives. If every lender in the market had cut by the same amount, β would be zero even though aggregate credit had fallen. What we recover is the reallocation of credit across lenders, which is the object of interest here, since the question is whether a disaster in Texas moves lending elsewhere rather than whether it lowers lending overall.

Second, borrowers may replace credit lost from an exposed lender by drawing more from an unexposed one. The borrower fixed effect nets that out by construction, so β is a differential and not the effect of exposure on a firm's total credit. How the two relate depends on how the firm's other lenders responded, which this design does not measure. If unexposed lenders expanded to fill the gap, the loss in total credit is smaller than β. If they cut as well, it is larger. Syndicated borrowers are the group best placed to replace a lender, so what the same shock would do to a firm with a single bank is not bounded by what we find here.

A weighting choice completes the specification. Estimated without weights, equation (10) gives every lender-borrower pair one vote, so β describes the typical lending relationship. Weighted by the pair's pre-storm loan amount, it gives every dollar one vote, so β describes the typical dollar of syndicated credit. The two coincide only when the response is unrelated to loan size, and the distinction matters here because the aggregate and stress-testing questions this paper asks are about dollars of credit rather than counts of relationships. Khwaja and Mian (2008) report both variants to verify that their conclusions are driven neither by the many small loans nor by a few very large ones, and we follow the same discipline. The intensive-margin estimates in the main text are dollar weighted, with the unweighted version of each exhibit in the Appendix. The extensive margin is unweighted throughout, since exit is a count of relationships and each relationship either survives or does not.

Equations (10) and (11) condition on the relationship surviving, since the change in lending exists only for pairs that lend in both windows. If exposed lenders are rebalancing away from hurricane risk, the sharper prediction is that they let relationships die. We therefore construct a quarterly panel of all pre-existing lender-borrower relationships and define an absorbing activity indicator, equal to one while the pair still originates and zero permanently after its last loan, and estimate

$$\mathbf{1}[\text{active}]_{ijt} = \gamma_{ij} + \alpha_{it} + \sum_{k \neq -1} \beta_k \, \mathbf{1}[t = k] \cdot \text{Harvey exposure}_j + \varepsilon_{ijt}, \tag{12}$$

with pair fixed effects γ_{ij} and borrower-quarter fixed effects α_{it}, so the comparison is again within borrower across lenders and demand is absorbed quarter by quarter as in equation (11). Each β_k measures the gap in the probability that the relationship is still active between a more and a less exposed lender of the same borrower. Because exit is a question about whether the relationship survives rather than about quantities, pairs are weighted

equally in this specification. Relationship counts, not dollars, are the natural metric for the extensive margin.

The assumption the design does require is that credit demand is common across a borrower's lenders, so that demand shocks load on α_i and do not vary systematically with the lender's Harvey exposure. It would fail if a borrower directed its post-Harvey demand toward particular lenders for reasons correlated with those lenders' Texas portfolios. A borrower contributes to the estimation only if it borrows from at least two lenders in both years, which the prevalence of syndication makes common.

The main threat to a supply interpretation is that borrowers themselves may be economically linked to the storm, through their own operations, supply chains, or regional demand. We address this in three ways. First, all borrowers in Harvey FEMA counties are dropped, and every result is reported alongside a version that drops borrowers in Hurricane Irma counties as well.[12] Second, we re-estimate the specification on samples that exclude all borrowers within 300 and 700 miles of the impact zone, distances at which any direct economic link to Harvey is implausible. Third, we re-run the entire design around placebo event dates in the third quarters of 2014, 2015, and 2016. The 2014 and 2015 windows end before any Harvey-related activity and are the clean falsification tests. The 2016 post window runs through September 2017 and so overlaps both the storm itself and the recovery of the oil market, which we use to read the energy-lending alternative rather than as a placebo.

The buffers do not absorb the hurricane-risk heterogeneity we examine next, because hurricane risk and distance to Harvey are correlated near the Gulf but not nationally. Historical hurricane risk runs up the entire Atlantic seaboard, far from Texas, so borrowers in high-risk counties remain plentiful even in the most distant samples. This is what allows the risk gradient to be separated from proximity to the storm.

4.4 Heterogeneity in the borrower's hurricane risk

Equation (10) asks whether exposed lenders cut credit. It does not ask which borrowers absorb the cut, and that is where the two competing readings of the result separate. If exposed lenders are simply short of funds after paying for the storm, the cut should fall on everyone alike. If instead they are rebalancing away from climate risk, the cut should sort

[12] Hurricane Irma made landfall in Florida on September 10, 2017, within the excluded quarter. Dropping Irma counties ensures the identification isolates spillovers to genuinely unaffected borrowers.

on the borrower's own exposure to hurricanes, regardless of how far that borrower sits from Texas. We therefore let the effect vary with the borrower's hurricane risk in two ways.

The first is a set of subsamples. We re-estimate equation (10) on samples that exclude all borrowers within 300 and 700 miles of the impact zone, and separately on samples that exclude Irma counties, financial borrowers, and each lender type in turn. Subsamples are transparent, but each one is estimated on different borrowers, so the coefficients are not directly comparable and precision falls as the sample shrinks.

The second is an interaction, which keeps every borrower in a single regression. We classify all non-Harvey borrowers by their county's historical hurricane count (Section 2.1), using the same ten-hurricane threshold that defines hurricane-prone counties throughout the paper, and estimate

$$\Delta \ln(\text{amount})_{ij} = \alpha_i + \beta\, \text{Harvey exposure}_j + \delta_H\, (\text{Harvey exposure}_j \times \text{High}_i) + \varepsilon_{ij}, \quad (13)$$

where Harvey exposure_j is standardized to unit variance and High_i marks borrowers in counties with ten or more historical hurricanes. Low-risk borrowers are the omitted reference, so β is the effect for a low-risk borrower and $\beta + \delta_H$ is the effect in the high-risk group. Each borrower belongs to exactly one group, so each group total is identified by that group's own borrowers and coincides with the corresponding subsample estimate, and the borrower fixed effect continues to absorb demand within every group. What the interaction adds is δ_H itself, the difference between the two groups estimated with a standard error, which the subsample splits cannot deliver. Appendix D reports the same regression with the borrowers split into thirds instead, following Gelman and Park (2009), which isolates the genuinely low-risk tail at the cost of smaller cells. These are the two coefficients that the framework of Section 3 maps to the two channels in equation (6). The rebalancing account predicts $\delta_H < 0$ and a group total near zero in the low-risk group. A funding shortfall predicts $\delta_H = 0$.

4.5 Market shares and portfolio shares

The within-borrower estimator identifies supply by holding the borrower fixed, which is also its limitation. It recovers how credit is reallocated across a borrower's lenders and is silent on whether the market as a whole moved. To ask that question we drop the borrower fixed effect and move to the lender level, collapsing the data to one observation per lender

parent and asking whether more exposed lenders shifted the geography of their lending. We measure two outcomes. The portfolio share is the fraction of a lender's own lending that goes to hurricane-risky counties, so it describes how the lender rearranged its own book. The market share is the lender's lending as a fraction of all lending in those counties, so it describes whether other lenders took the space that was vacated. For each outcome we regress the post-minus-pre change on the lender's Harvey exposure,

$$\Delta\text{share}_j = \alpha + \beta \cdot \text{Harvey exposure}_j + \varepsilon_j \tag{14}$$

estimated by weighted least squares with pre-period lending as the weight, so each lender counts in proportion to the credit it supplied and the estimates describe the reallocation of actual credit rather than of lender counts.[13] The two exercises are complementary. The within-borrower design gives the causal estimate, and the share regressions say whether it is large enough to show up in the aggregate.

5 Results and Discussion

In this section, we first document what happened inside the disaster zone, where borrowing expanded and every lender took part in the expansion. We then turn to the unaffected borrowers. On the intensive margin we report the credit supply contraction, its dynamics, its insensitivity to distance from the storm, and its concentration among hurricane-risky borrowers. The extensive margin then follows the exit of lending relationships. Finally, a lender-level analysis of portfolio and market shares shows the aggregate footprint of the reallocation. We close with the robustness checks and the interpretation of the evidence.

5.1 The credit market inside the disaster zone

Table 1 reports equation (8). It asks whether the borrowers whose counties Harvey damaged ended the following year with more syndicated credit or with less, and it asks that question twice, once with damage as an indicator for any recorded loss and once with the

[13] The largest parent carries 8.8 percent of the regression weight and the top ten carry 55.5 percent, so the weighted estimates describe the reallocation of actual credit volumes. Unweighted estimates are smaller and noisier, indicating that the reallocation is concentrated among large lenders.

dollar value, so that the answer does not rest on how damage is coded.

Table 1: Borrower Credit Growth and Harvey Damage in the Disaster Zone

	$\Delta \ln$(total credit)			
	(1)	(2)	(3)	(4)
$\mathbf{1}[\text{Damage} > 0]$	0.121***	0.125***	0.126***	0.176*
	(0.040)	(0.034)	(0.033)	(0.091)
$\ln(1 + \text{Damage})$	0.0058***	0.0057***	0.0058***	0.0081**
	(0.0016)	(0.0015)	(0.0014)	(0.0039)
N (borrowers)	917	917	916	910
Borrower controls	No	Yes	Yes	Yes
Industry FE	No	No	Yes	Yes
State FE	No	No	No	Yes

Notes: Cross-sectional regressions of borrower credit growth on Harvey damage in the borrower's headquarters county, equation (8), with each cell a separate regression. The pre-period is July 2016 through June 2017 and the post-period October 2017 through September 2018; the hurricane quarter (2017Q3) is excluded and Irma-only borrowers are dropped. Damage is the sum of flooding and hurricane damage recorded by SHELDUS in August and September 2017 in the counties with FEMA major disaster declarations for Harvey, and zero elsewhere. The outcome is the log change in total credit for borrowers active in both windows. Borrower controls are the log of pre-period credit and the number of pre-period lenders. Standard errors clustered by county in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Borrowers headquartered in damaged counties increased their total syndicated borrowing by about 12 log points relative to other borrowers, and the estimate barely moves as controls are added. The dollar value of damage matters beyond the indicator. Within the set of damaged-county borrowers, credit growth rises with the intensity of damage in the borrower's county. The disaster zone did not lose credit after Harvey. It gained credit, in line with the rebuilding demand that Cortés (2014) documents after disasters and with the direct effect that Ivanov et al. (2022) estimate across disaster episodes.

The natural follow-up is who financed this surge. One reading of the contraction we document later is that exposed lenders lost lending capacity, absorbing losses on their Harvey-county book and so having fewer dollars to lend. If that is what happened, those lenders should have ceded the disaster-zone expansion to their unexposed competitors, because a

lender short of capacity cannot fund a surge in demand. Table 2 re-estimates equation (10), dollar weighted as in Table 3, separately on the borrower groups that the main analysis excludes.

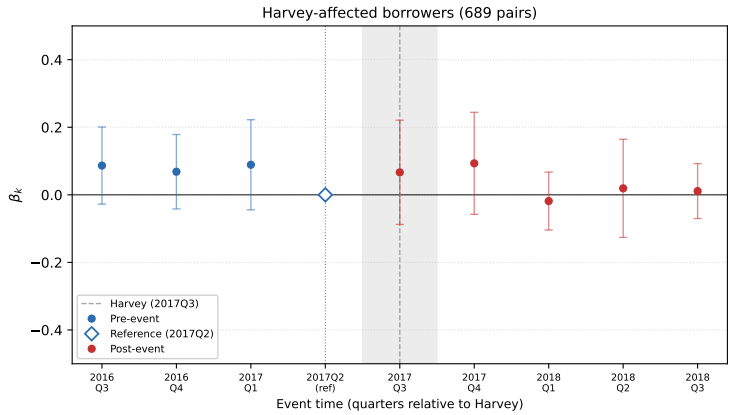
Table 2: Khwaja–Mian Estimates by Borrower Group

Outcome: $\Delta \ln(amount)$	Harvey-affected
Harvey exposure (SD), dollar weighted	0.013
	(0.038)
Harvey exposure (SD), unweighted	0.001
	(0.021)
Borrower FE	Yes
N (pairs)	689
Lender parents	99

Notes: Each row is a separate estimate of equation (10) on Harvey-affected borrowers, the group the main analysis sets aside. The first row weights each pair by its pre-storm loan amount and the second weights pairs equally, as in Table 3. In a single interacted regression that pools these borrowers with the unaffected borrowers of column (2) of Table 3, the Harvey-affected coefficient exceeds the unaffected one by 0.053 dollar weighted and 0.026 unweighted ($p = 0.23$ in both). Borrower fixed effects included; standard errors clustered by lender parent in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

For Harvey-affected borrowers the within-borrower coefficient is a statistical zero in both weightings. A lender one standard deviation more exposed to Harvey expanded credit to a damaged-county borrower at the same rate as that borrower's other lenders. Because the borrower fixed effect absorbs the common component of each borrower's credit growth, this zero sits on top of the credit expansion in Table 1. Every lender, exposed or not, participated in the disaster-zone expansion in proportion to its existing relationship. This is what the disaster-zone literature would lead one to expect, since established lending relationships shield affected borrowers' access to recovery credit (Berg and Schrader, 2012) and lenders with local exposure lean into recovery lending rather than away from it (Cortés, 2014; Koetter et al., 2020). Figure 2 confirms that this equality holds quarter by quarter, with no pre-trend and no post-storm divergence.

Figure 2: Event Study. Harvey-Affected Borrowers



Notes: Event-study estimates of equation (11) on Harvey-affected borrowers only (689 lender-borrower pairs). The outcome is log lending volume at the lender-borrower-quarter level; borrower-quarter and lender-parent fixed effects included; the omitted period is 2017Q2. The joint test that the pre-event coefficients are zero has $p = 0.48$; no individual coefficient is significant at the ten percent level. Vertical bars denote 95% confidence intervals; standard errors clustered by lender parent.

Market outcomes in the disaster zone tell the same story. Exposed lenders grew there in step with the market, neither gaining nor ceding Harvey-county market share, which we show with the lender-level design of Section 4.5 when we turn to it in Section 5.4. Borrowers in Irma counties sharpen the reading. The regressor there is still the lender's Harvey exposure, not its Irma exposure, and the borrower fixed effect absorbs whatever Irma did to the borrower's own demand. Where Harvey-exposed lenders had no rebuilding demand of their own to meet, the same specification yields a cut of 9.9 percent per standard deviation of Harvey exposure (s.e. 3.8, on 454 pairs at 96 lender parents). The zero for Harvey-affected borrowers is therefore not a general inability to adjust. It reflects rebuilding demand that every lender chose to accommodate. The reallocation this paper documents happened outside the disaster zone, and the rest of this section traces it there.

5.2 The intensive margin in unaffected areas

Table 3 reports the pooled estimate of equation (10), with each pair weighted by its pre-storm loan amount and the unweighted estimate beneath it. Column (1) drops Harvey-

county borrowers and column (2) additionally drops Irma-county borrowers.[14] A borrower's credit from a lender one standard deviation more exposed to Harvey falls by 4.8 percent relative to its other lenders, significant at the one percent level. Additionally dropping Irma-county borrowers reduces the estimate to 4.0 percent, significant at the five percent level on the smaller sample. Counting relationships instead of dollars, the contraction is 3.4 and 2.4 percent. The dollar-weighted response is the larger of the two in both samples, so the cuts fell disproportionately on the larger loans. The magnitude is close to the roughly 3 percent contraction that Ivanov et al. (2022) estimate for borrowers in distant unaffected counties using within-county variation. Recovering a similar figure one level deeper, within the borrower rather than within its county, suggests that the spillover they document is indeed a supply response by exposed lenders rather than a residual demand difference across borrowers in the same county.

Table 3: Khwaja–Mian Estimates. Pooled Sample

	$\Delta \ln(amount)$	
	(1) Drop Harvey	(2) Drop Harvey+Irma
Harvey exposure (SD), dollar weighted	-0.048^{***}	-0.040^{**}
	(0.017)	(0.019)
Harvey exposure (SD), unweighted	-0.034^{***}	-0.024^{*}
	(0.013)	(0.014)
Borrower FE	Yes	Yes
N (pairs)	4,727	4,273
Lender parents	242	233

Notes: Each observation is a lender parent × borrower pair present in both the pre-period (July 2016 through June 2017) and the post-period (October 2017 through September 2018); the hurricane quarter (2017Q3) is excluded. Harvey exposure is the lender's pre-storm portfolio share in Harvey FEMA counties, standardized to unit variance. Column (2) additionally drops borrowers in Hurricane Irma counties. Each row pair is a separate regression: the dollar-weighted rows weight each pair by its pre-storm loan amount, the unweighted rows give each pair equal weight (Section 4.3). Borrower fixed effects included; standard errors clustered by lender parent in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

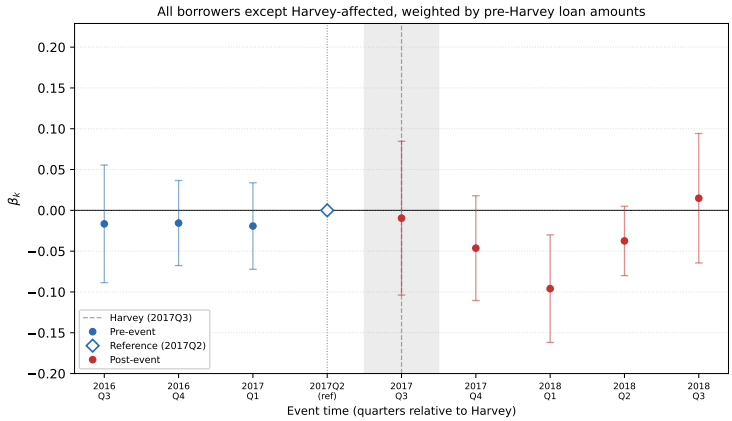
[14]Hurricane Irma made landfall in Florida on September 10, 2017, inside the excluded quarter, so a borrower in an Irma county experienced a storm of its own during the event window. Dropping these borrowers, as set out in Section 4.3, leaves a sample in which no borrower was touched by either storm.

The adjustment is entirely in quantities. The same specification estimated on the change in the loan spread and in maturity yields coefficients that are small and statistically indistinguishable from zero in every sample we consider.[15] Exposed lenders did not reprice credit to unaffected borrowers. They supplied less of it. This null needs to be read against the pricing evidence with some care, because the two designs do not ask the same question. Correa et al. (2026) and Javadi and Masum (2021) both hold the lender fixed and place the variation on the borrower, so they recover the repricing that is common to all of a borrower's lenders. Our borrower fixed effect absorbs exactly that common movement. A spread increase applied by every lender at once is invisible to equation (10) by construction, so finding nothing here is not evidence against their result. What this design can see is a spread that the exposed lenders charge over and above the borrower's other lenders, and there is none. Their reassessment shows up in how much they are willing to lend rather than in what they charge relative to their peers, which is what one would expect in a market where the spread is set jointly for the syndicate and participation is not. We return to this distinction in Section 5.6.

The pooled estimate compresses the adjustment into a single pre versus post comparison. Figure 3 traces it quarter by quarter, reporting the dynamic specification of equation (11), estimated on all non-Harvey-affected borrowers, dollar weighted with each pair weighted by its mean pre-period quarterly loan amount.

[15] On the sample of column (1), the spread analog of equation (10) yields −0.11 basis points per standard deviation of exposure (s.e. 0.48) unweighted and 0.81 basis points (s.e. 1.05) dollar weighted. The maturity analog yields −0.03 months (s.e. 0.10) unweighted and −0.08 months (s.e. 0.19) dollar weighted. A tranche-level test of origination spreads on lender exposure gives −3.5 basis points (s.e. 8.5). None approaches significance.

Figure 3: Event Study. All Non-Harvey-Affected Borrowers



Notes: Event-study estimates of equation (11) on all non-Harvey-affected borrowers, with each pair weighted by its mean pre-period quarterly loan amount. The outcome is log lending volume at the lender-borrower-quarter level; borrower-quarter and lender fixed effects included; the omitted period is 2017Q2. The joint test of the pre-event coefficients does not reject ($p = 0.90$). Vertical bars denote 95% confidence intervals; standard errors clustered by lender parent.

The pre-event coefficients sit on zero. After landfall the contraction in dollars builds over two quarters to a trough of 9.6 percent per standard deviation of exposure in the first quarter of 2018 before flows recover by the third quarter of that year. The gradual build-up is inconsistent with a mechanical liquidity shock at landfall and closer to portfolio rebalancing over the following year, a timing that also fits the salience channel, in which the reassessment spreads as the event dominates attention rather than arriving on the day of landfall (Dessaint and Matray, 2017; Correa et al., 2026). The recovery in flows does not mean the adjustment was undone. As Section 5.3 shows, the extensive margin deteriorates precisely when flows recover. Dropping Irma-county borrowers as well, as in column (2) of Table 3, leaves the path unchanged. The pre-trends stay flat and the trough is slightly deeper than in the full sample (Figure A3 in Appendix E). The dynamics are not driven by Irma-hit borrowers demanding credit at the same time.

Borrowers close to the impact zone could be affected by Harvey through their own operations, supply chains, or local demand, which would contaminate a pure credit-supply reading. Table 4 re-estimates equation (10), dollar weighted as in Table 3, on progressively stricter samples that exclude all borrowers within 300 and 700 miles of the Harvey impact zone (cumulative exclusion buffers, not mutually exclusive rings).

Table 4: Khwaja–Mian Estimates Excluding Borrowers Near the Impact Zone

Outcome: $\Delta \ln(amount)$	Pooled	>300 mi	>700 mi
Panel A: Drop Harvey borrowers (main)			
Harvey exposure (SD)	-0.048^{***}	-0.047^{***}	-0.043^{**}
	(0.017)	(0.018)	(0.019)
N	4,727	4,493	3,809
Panel B: Drop Harvey and Irma borrowers			
Harvey + Irma exposure (SD)	-0.043^{**}	-0.045^{**}	-0.044^{*}
	(0.020)	(0.021)	(0.023)
N	4,273	4,039	3,545
Borrower FE	Yes	Yes	Yes

Notes: Each column beyond the first restricts the sample to borrowers located more than 300 or 700 miles from the nearest Harvey-affected county. The buffers are cumulative exclusions, not mutually exclusive rings. In Panel B the exposure measure is the lender's pre-storm portfolio share in Harvey and Irma counties combined (mean 21.3 percent of the portfolio, against 13.4 percent for Harvey alone), standardized on the full sample, so the shock is the lender's exposure to the entire 2017 hurricane season. Each pair is weighted by its pre-storm loan amount, as in the headline rows of Table 3; the unweighted version is Table A7 in the Appendix. Borrower fixed effects included; standard errors clustered by lender parent in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The contraction does not come from borrowers close to the storm. The dollar-weighted estimate barely moves across the buffers, from 4.8 percent pooled to 4.7 and then 4.3 percent as everyone within 300 and then 700 miles is excluded. Panel B widens the shock from Harvey to the entire 2017 hurricane season, measuring each lender's pre-storm portfolio share in Harvey and Irma counties combined, and the picture is unchanged, between 4.3 and 4.5 percent across the buffers. The unweighted version, Table A7 in Appendix D, shows the same pattern. Because each buffer removes any borrower who could plausibly be economically linked to the storm, the surviving effect must operate through the lender side. The insensitivity to distance is itself informative. Proximity to the storm proxies for a borrower's own economic link to Harvey, so a contraction that is as strong 700 miles away as it is nearby is hard to attribute to anything about the borrower. Geography is evidently not the margin on which lenders acted, which motivates asking what is.

Distance from the impact zone is an identification device, not a mechanism. If exposed lenders are rebalancing away from climate risk, the supply cut should sort on the borrower's

own hurricane exposure, not on miles from Texas. We look at this in two ways, first by splitting borrowers at a series of risk thresholds and then by letting the effect vary above and below the ten-hurricane threshold inside a single regression.

Table 5 takes the first route. It re-estimates equation (10), dollar weighted as in Table 3, separately on low- and high-risk borrowers, cutting the county hurricane count at three, five, and ten historical passages. The three cuts are not chosen to be optimal, they are reported together to show that the result does not depend on where the line is drawn.

Table 5: Supply Contraction by Borrower Hurricane-Risk Threshold

	Low-risk borrowers			High-risk borrowers		
Outcome: $\Delta \ln(amount)$	<3	<5	<10	≥ 3	≥ 5	≥ 10
Harvey exposure (SD)	−0.046	−0.030	−0.011	−0.048***	−0.054***	−0.074***
	(0.038)	(0.032)	(0.023)	(0.018)	(0.018)	(0.022)
Borrower FE	Yes	Yes	Yes	Yes	Yes	Yes
N (pairs)	784	1,173	1,871	3,943	3,554	2,856
Lender parents	100	133	167	217	205	182

Notes: Each entry is a separate estimate of equation (10) on the indicated subsample of non-Harvey borrowers, with each pair weighted by its pre-storm loan amount as in the headline rows of Table 3; the unweighted version is Table A8 in the Appendix. Borrower hurricane risk is the number of hurricanes passing within 100 miles of the borrower's headquarters county between 1851 and Harvey, from NOAA's HURDAT2 storm-track database. Thresholds are cumulative, so the low- and high-risk columns at a given cut partition the sample. Borrower fixed effects included; standard errors clustered by lender parent in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Among borrowers whose county has seen at least ten hurricanes, the supply cut is 7.4 percent of loan volume per standard deviation of exposure, roughly 50 percent larger than the pooled estimate, and the high-risk cell is significant at the one percent level at every threshold, growing as the cut rises. Among low-risk borrowers the estimate is never statistically distinguishable from zero at any cut. The unweighted version (Table A8) shows the same pattern with a low-risk side that is flat between 1.3 and 1.8 percent at every cut. This is the pattern a portfolio-rebalancing account predicts and the one a funding shortfall does not, since a lender short of funds has no reason to concentrate the cut on borrowers whose only distinguishing feature is a hurricane history of their own. The sorting on the borrower's own disaster risk echoes Rehbein and Ongena (2022), who find that banks hit by a flood

retrench most from regions they anticipate as future disaster risks, which suggests that this forward-looking margin of adjustment generalizes across shocks and lending markets.

Table 6 takes the second route, letting the effect vary above and below the ten-hurricane threshold inside a single regression, dollar weighted as in Table 3. The cut is not chosen after the fact. Ten historical hurricanes is the definition of a hurricane-prone county used throughout the paper, and the group totals coincide by construction with the ten-hurricane columns of Table 5. What the single regression adds is the difference between the two groups, δ_H, estimated with a standard error.

Table 6: Supply Contraction by the Borrower's Own Hurricane Risk

Outcome: $\Delta \ln(amount)$	Low risk, <10 (base)	High risk, ≥10	Difference δ_H
Harvey exposure (SD)	−0.011	−0.074***	−0.064**
	(0.023)	(0.022)	(0.030)
N (pairs)	1,871	2,856	

Notes: One regression on all non-Harvey borrowers, equation (13), with borrowers split at ten historical hurricanes in their county (hurricanes within 100 miles of the county centroid, 1851 to Harvey, from NOAA HURDAT2), the threshold that defines hurricane-prone counties throughout the paper. Each pair is weighted by its pre-storm loan amount as in the headline rows of Table 3; the unweighted version is Table A9 in the Appendix and the finer three-band cut is Table A10 in the Appendix. The low-risk group is the omitted reference; the high-risk total is computed by linear combination and coincides with the ten-hurricane column of Table 5. The three cells come from one regression on 4,727 pairs; the N row reports the pairs in each group. Borrower fixed effects included; standard errors clustered by lender parent in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The difference is large and statistically significant. Low-risk borrowers see essentially no cut, a point estimate of 1.1 percent with a tight standard error, while high-risk borrowers lose 7.4 percent of loan volume per standard deviation of exposure, and the gap between the two is significant at the five percent level. This is the formal version of the pattern in Table 5. Not only is the high-risk cell significant and the low-risk cell not, the two are significantly different from each other. The unweighted regression, Table A9 in Appendix D, shows the same gradient with a smaller and marginally significant gap. At the looser cuts of three and five hurricanes the difference is not significant, consistent with Table 5, where the gradient sharpens as the low-risk group becomes genuinely low risk. The finer three-band cut, Table A10 in Appendix D, tells the same story. Borrowers in counties with at most 6

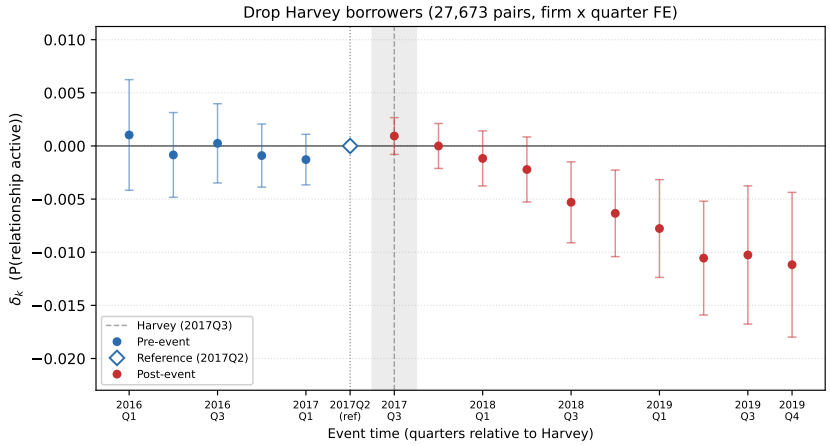
historical hurricanes are indistinguishable from zero, while borrowers with 7 or more carry the cut, 5 to 6 percent per standard deviation of exposure. The dynamics agree. An event study estimated separately for the two risk groups shows flat pre-trends and a deep early-2018 trough for high-risk borrowers, while the low-risk path stays shallow after the storm quarter, as shown in Figure A2 in Appendix D. The risk gradient is the same borrower margin on which Correa et al. (2026) find spreads moving after climate-linked disasters, which suggests, without proving, that quantity and price adjustments after salient storms are two expressions of the same reassessment of the borrower's physical risk.

5.3 The extensive margin in unaffected areas

The intensive margin measures how much surviving relationships shrink. Equation (12) asks whether exposed lenders let relationships die altogether, following an absorbing activity indicator that equals one while the pair still originates and zero permanently after its last loan, with pair and borrower-quarter fixed effects so the comparison stays within borrower across lenders.

Figure 4 shows the result. Pre-event coefficients are flat. Nothing happens in the first year after the storm, the period when amounts are being cut. From the third quarter of 2018 the probability that a relationship remains active falls, reaching 1.1 percentage points per standard deviation of exposure by the end of 2019. Because exit is a count of relationships, equal weights are the natural metric here. Exposed lenders let relationships lapse two to four quarters after the amount cuts. The sequence, amounts first and terminations later, is what relationship rents predict. A lender that values a continuing relationship absorbs a small desired reduction inside it and severs it only when the desired cut is large enough to exhaust the rent, so exit is the margin of last resort and should lag the reduction in amounts. Severed credit relationships are the channel through which lender distress reached employment in Chodorow-Reich (2014), and the timing here suggests lenders shaved amounts before letting relationships die.

Figure 4: Exit Event Study. All Non-Harvey-Affected Borrowers

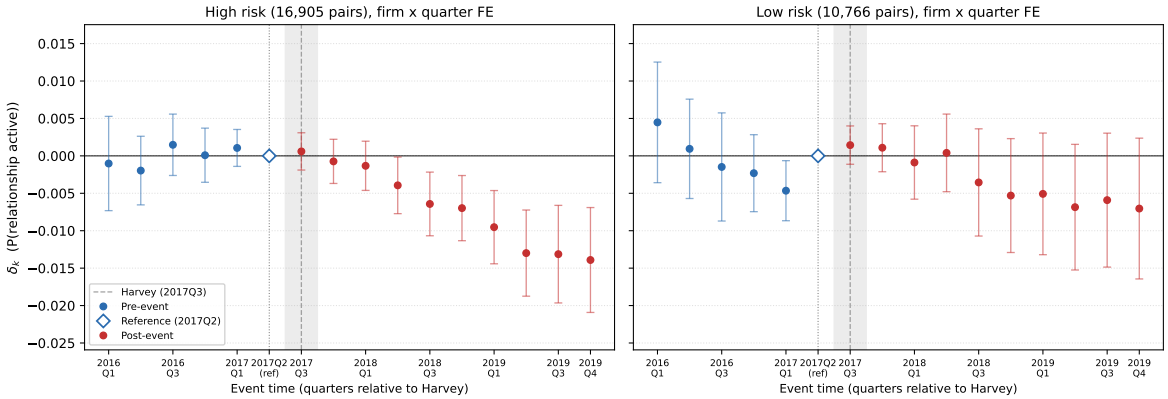


Notes: Event-study estimates of equation (12). The outcome equals one while the lender-borrower pair still originates and zero permanently after its last loan. Each pair is weighted equally: exit is a question about whether the relationship survives, not about quantities, so relationship counts are the natural metric. Pair and borrower-quarter fixed effects included; the omitted period is 2017Q2. Vertical bars denote 95% confidence intervals; standard errors clustered by lender parent. The joint test of the pre-event coefficients does not reject ($p = 0.65$).

Figure 5 splits the exit event study at 10 historical hurricanes within 100 miles of the borrower's county, the same threshold that separates high- from low-risk borrowers in Table 6, mapped in Figure 1. The lapse is concentrated in high-risk borrowers, reaching 1.4 percentage points per standard deviation of exposure by the end of 2019, while low-risk borrowers show an insignificant decline of half that size.[16] The intensive-margin gradient in amounts reappears in who loses the relationship. The borrowers that exposed lenders stop lending to are the hurricane-risky ones.

[16]The high-risk panel shows clean pre-trends. The low-risk panel shows some pre-event instability, so we read the low-risk dynamics with caution. The conclusion that the exit is concentrated among high-risk borrowers rests on the high-risk panel.

Figure 5: Exit by Borrower Hurricane Risk



Notes: Same specification as Figure 4, splitting borrowers at 10 historical hurricanes within 100 miles of the borrower's county (high risk ≥10, low risk <10). Each pair is weighted equally. Vertical bars denote 95% confidence intervals; standard errors clustered by lender parent.

5.4 Market shares and reallocation across risky and safe counties

Both margins say that exposed lenders pulled back from hurricane-risky borrowers. If that is what happened, the reallocation should also be visible in market outcomes. We collapse the data to one observation per lender parent, using the same pre- and post-windows as Table 3, and define risky counties as those with at least ten historical hurricane passages. No borrowers are excluded. The Harvey and Irma counties, all of which meet the risky definition, stay in both the numerator and the denominator, and the table separates their contribution. We then estimate equation (14) by weighted least squares, as set out in Section 4.5.

Table 7: Market Shares. Reallocation Across Risky and Safe Counties

	Δ Portfolio share		Δ Market share	
Sample	Risky counties	Safe counties	Risky counties	Safe counties
All lenders	-1.86^{***}	1.86^{***}	-0.165^{***}	0.004
	(0.49)	(0.49)	(0.054)	(0.033)
of which: outside the disaster areas	−0.80		-0.174^{**}	
	(0.54)		(0.068)	
of which: Harvey and Irma counties	-1.06^{*}		-0.129^{*}	
	(0.54)		(0.066)	
Lender parents	357	357	357	357

Notes: Each entry is the coefficient on the lender's standardized Harvey exposure, so it is the change in the outcome associated with a one standard deviation increase in exposure. There is one observation per lender parent, the windows are those of Table 3, and no borrowers are excluded. Risky counties have at least 10 historical hurricane passages within 100 miles; every Harvey and Irma county meets this definition, so the two indented rows partition risky-county lending and their portfolio-share coefficients sum to the first row. Both outcomes are post-minus-pre changes measured in percentage points. Portfolio share is the lender's own lending in risky counties as a fraction of its total lending, and market share is the lender's lending as a fraction of all lending in those counties. Because risky and safe counties partition each lender's portfolio, the safe-county portfolio-share column is the mirror image of the risky one by construction. The mean pre-period portfolio share in risky counties is 65.8 percent. Within the disaster areas, the market-share decline comes entirely from the Irma counties (-0.32, $p = 0.01$). The Harvey-county market share is unchanged (-0.02, $p = 0.77$), which is the fact used in Section 5.1. Estimated by weighted least squares with pre-period lending volume as the weight. Every entry is a separate regression. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 7 shows the reallocation. A one standard deviation more exposed lender moved about 1.9 percentage points of its portfolio out of risky counties, off a pre-storm base of 65.8 percent. In market-share terms, exposed lenders ceded share only in risky counties, while in safe counties their share is unchanged. The decomposition rows show where the retreat happened. The share ceded in risky counties comes from counties outside the disaster areas and from the Irma counties, not from the Harvey counties, where the borrowing surge of Section 5.1 was concentrated and where exposed lenders held their market share exactly. The safe-county market-share column carries its own message. Exposed lenders' market share in safe counties did not rise, so the portfolio tilt toward safe counties reflects cutting risky lending, which shrinks the denominator, rather than relocating those dollars. This is a contraction of risky credit, not an expansion elsewhere. The result is not an artifact of disaster-zone flows, since excluding the Harvey and Irma counties from both the numerator

and the denominator leaves the portfolio shift and the risky market-share loss slightly larger. The within-borrower supply cut of Table 3 therefore aggregates to a visible shift of the lending market in hurricane-risky counties toward less-exposed lenders. The geography of this reallocation is telling. In Cortés and Strahan (2017) the map of the adjustment is drawn by the disaster's demand, with funds flowing toward the damaged region and away from wherever they can be spared. Here the map is drawn by the destination counties' hurricane risk, which is the footprint one would expect if exposed lenders were managing climate exposure rather than funding needs.

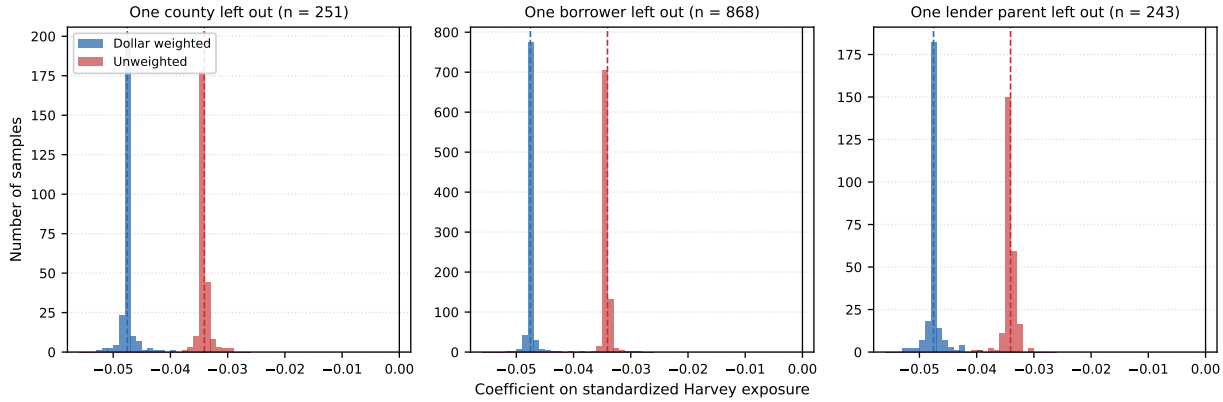
5.5 Robustness

This section probes the result, ordered from the checks that discipline the headline estimate most directly to the more specialized variations. We first show that no single county, borrower, or lender parent drives the contraction. We then remove financial borrowers, who could respond to the storm through channels of their own, and split the sample by lender type. Next we rebuild the two constructed inputs, the lending shares under alternative imputation rules and the borrower's hurricane risk under a recency-weighted measure. We close with placebo event dates, which re-run the entire design in years without a comparable storm.

5.5.1 Leave-one-out samples

If the estimate rested on a single influential county, borrower, or lender, dropping that unit would move it visibly. Figure 6 re-estimates the pooled specification of Table 3, column (1), dropping one borrower county, one borrower, or one lender parent at a time, and plots the distribution of the resulting coefficients in each weighting. All 2,724 re-estimates are negative and significant at the five percent level. The most influential single unit moves the dollar-weighted estimate from -0.048 to -0.040, about half a standard error, and the unweighted estimate from -0.034 to -0.029. No single county, borrower, or lender parent drives the contraction.

Figure 6: Leave-One-Out Distributions of the Pooled Estimate



Notes: Histograms of the coefficient on standardized Harvey exposure from re-estimating the pooled specification of Table 3, column (1), dropping one borrower county (251 samples), one borrower (868), or one lender parent (243) at a time, with exposure standardized once on the full sample. Blue: each pair weighted by its pre-storm loan amount; red: equal weights. Dashed lines mark the full-sample baselines (−0.048 dollar weighted, −0.034 unweighted) and the solid line marks zero. Every re-estimate is negative and significant at the five percent level. Borrower fixed effects included; standard errors clustered by lender parent.

5.5.2 Excluding financial borrowers

About 17 percent of borrowers are financial firms, which could respond to Harvey through channels of their own, such as insurance losses.[17] Dropping them and re-estimating the pooled specification yields a contraction of 3.7 percent per standard deviation of exposure, slightly larger than the full-sample estimate. The event study on this sample, weighted by pre-storm loan amounts as in the headline series, shows the same picture, flat pre-trends and a trough of 9.4 percent in early 2018 (Figure A4 in Appendix E). The results are not driven by financial-sector borrowers.

5.5.3 Lender types

Table 8 re-estimates equation (10) separately for U.S. banks, foreign banks, and nonbanks, with exposure standardized once on the full sample so the columns are comparable.

[17] Financial borrowers are those whose modal SIC code falls in 6000 to 6999, which is broader than the financial services group of Table A3 because it also takes in real estate and REITs.

Table 8: Subsamples by Lender Type

Outcome: $\Delta \ln(amount)$	US bank	Foreign bank	Nonbank/other
Drop Harvey	-0.044^{***}	-0.023	-0.119
	(0.015)	(0.018)	(0.113)
Drop Harvey + Irma	-0.041^{**}	-0.018	-0.053
	(0.017)	(0.018)	(0.065)
Pairs (drop Harvey)	2,520	1,333	320
Lender parents	88	71	37

Notes: Equation (10) estimated separately by lender type, with exposure standardized once on the full sample. Lender type from DealScan's lender classification; a parent takes the type accounting for most of its entity appearances. Borrower fixed effects are imposed within type, so a borrower contributes only if it has at least two lenders of that type. Standard errors clustered by lender parent in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The contraction is precisely estimated only for U.S. banks, at 4.1 to 4.4 percent per standard deviation of exposure. Foreign banks cut by about half as much and the estimate is not significant. The nonbank point estimate is the largest of the three but rests on few pairs and parents, and its standard error is wide enough that nothing can be concluded from it. A pooled interaction cannot reject equality of the effect across types, so the subsamples rank the channel without sharply separating it. That the precise estimate sits with U.S. banks is natural, since they hold most of the Harvey-county credit and face the domestic supervisory climate agenda most directly, though the wide intervals for the other types mean we cannot rule out similar behavior there.

5.5.4 Alternative lending-share imputations

Section 2.2 describes the baseline rule for the syndicate shares, under which reported shares are always kept and missing shares receive the Chodorow-Reich (2014) role-by-structure averages. Table A11 in Appendix E re-estimates the pooled specification of Table 3, column (1), keeping the reported shares and imputing the missing ones under two alternative rules, an equal split among the non-reporting lenders of each tranche, and a split proportional to each non-reporting lender's size in the pre-storm syndicated loan market. Lender exposure and loan amounts are rebuilt under each rule. The specification, sample, and standardization

are unchanged, and the re-estimates are unweighted, so the baseline they should be compared with is 3.4 percent. The estimate is −3.0 percent per standard deviation of exposure under the equal split, significant at the five percent level, and −2.5 percent under the size-based split, similar in magnitude but less precisely estimated. The contraction does not hinge on the imputation rule.

5.5.5 A recency-weighted measure of hurricane risk

The hurricane count of Section 2.1 treats a storm in 1851 and a storm in 2016 as equally informative about a county's risk. If lenders read the recent record rather than the full history, that measure is mismeasured in a way that could either mask or manufacture the risk gradient of Table A10. We therefore rebuild the borrower's hurricane risk as an exponentially decayed count, in which a storm passing within 100 miles of the county centroid contributes weight $0.5^{(\text{years before Harvey})/h}$ rather than one, for half-lives h of 5, 10, 20, and 30 years. A storm thirty years before Harvey enters with weight 1.6 percent under a five-year half-life and 50 percent under a thirty-year half-life, so the four measures span a wide range of assumptions about the memory lenders bring to the problem. We re-cut terciles on each decayed measure and re-estimate the saturated three-group regression of equation (A17), with the groups defined by terciles of the decayed count. Because each borrower belongs to exactly one tercile and borrower fixed effects are absorbed within borrower, each group total is identified only by the borrowers in that group.

Table A12 reports the results. The gradient is stable across every half-life. The upper third contracts by 5.3 to 6.1 percent per standard deviation of exposure, significant at the five percent level or better throughout, while the lower third stays between 1.3 and 2.3 percent and is never distinguishable from zero. Discounting old storms neither creates nor destroys the pattern of Table A10, so the risk gradient does not rest on how far back the hurricane record is allowed to reach.

The two aggregate exercises were also repeated with the decayed measure, splitting counties at the median of the five-year half-life rather than at ten historical passages. The market-share reallocation of Table 7 survives. A one standard deviation more exposed lender moved 1.8 percentage points of its portfolio out of risky counties, and ceded market share in risky counties but not in safe ones. The exit gradient of Figure 5 does not survive as cleanly. Under the median split both groups exit by the end of 2019, at 1.3 percentage points per standard

deviation of exposure for high-risk borrowers and 1.0 for low-risk borrowers, counting relationships equally as in Figure 5. The median cut of the decayed measure places roughly half of lending in the risky group, against the 56.5 percent portfolio share of the ten-passage cut in the same sample, which excludes the affected counties, and assigns counties with modest recent activity to the low-risk group, so the two groups are closer together than under the baseline split. The extensive-margin gradient therefore rests on the ten-passage definition, while the level of the exit response does not.

5.5.6 Placebo event dates

If Texas-concentrated lenders were simply on different lending trends, the design would find effects at any date. Table A13 in Appendix E repeats the entire design, the same Harvey counties, exposure measure, windows, distance buffers, and dollar weighting, around placebo event dates in the third quarters of 2014, 2015, and 2016. The 2014 and 2015 windows end before any Harvey-related activity. The 2016 post window runs through September 2017, so it overlaps both the quarter of Harvey's landfall and the recovery of the oil market, and we discuss it separately below.

Neither clean placebo date produces a contraction. Every coefficient for 2014 and 2015 is small and statistically indistinguishable from zero in every distance sample and in both weightings, with the unweighted version in Table A14. The 2016 placebo is informative for a different reason. Its windows straddle the turnaround of the oil market. The pre window, July 2015 through June 2016, covers the trough of the 2014 to 2016 oil price collapse, when banks were writing down syndicated energy loans on a large scale, and the post window opens with OPEC's production cut of November 2016 and the price recovery that followed.[18] Because the Harvey counties sit at the center of the U.S. oil and gas industry, lenders concentrated there are plausibly also concentrated in energy lending, so an oil recovery predicts an expansion by exposed lenders in the 2016 windows. That is what the table shows. Dollar weighted, exposed lenders expand by 3.4 to 3.7 percent, the opposite sign of a supply cut, while the unweighted estimate is flat, consistent with a rebound running

[18] West Texas Intermediate fell from above $100 per barrel in mid-2014 to a low of $26 in February 2016, inside the placebo pre window. On November 30, 2016, at the start of the placebo post window, OPEC agreed to its first production cut since 2008 and prices recovered into the $50s through 2017. In the interagency Shared National Credit reviews of syndicated loans, classified oil and gas commitments rose from $6.9 billion in 2014 to $77 billion in early 2016, 27 percent of all classified commitments.

through large energy facilities. The dynamic version, Figure A5 in Appendix E, confirms the timing. Traced quarter by quarter, the weighted 2016 path has no significant quarter before the OPEC cut, builds to an expansion by the second quarter of 2017 as prices stabilized, and is flat in the quarter of Harvey's landfall itself. Pre-storm concentration in the Harvey counties predicts lending cuts around Harvey and lending growth when oil recovers. The contraction is a response to the event, not a stable trait of Texas-concentrated lenders.

5.6 Interpretation as climate-risk reallocation

Section 3 reduced the interpretation to a single coefficient. A capacity shock alone implies $\delta_H = 0$, a uniform contraction across borrowers. Belief updating implies $\delta_H < 0$ with a low-risk group total near zero. The estimates deliver the second pattern. Four further features of the evidence discipline the reading. First, exposed lenders cut credit to unaffected borrowers in quantities, not spreads. Second, what matters is not the borrower's distance from Harvey but the borrower's own hurricane risk. The contraction is flat in distance buffers out to 700 miles, yet sorts on the county's hurricane history and shows up in market shares only in hurricane-risky counties. Third, the adjustment follows a sequence. Amounts are cut in the first three quarters, and relationships with high-risk borrowers are allowed to lapse two to four quarters later. Fourth, the placebo dates show no contraction, so the behavior is a response to Harvey rather than a trait of Texas-concentrated lenders.

A pure balance-sheet impairment channel fits this pattern poorly. Lenders facing losses cut credit broadly, wherever they can (Peek and Rosengren, 2000), whereas the contraction we document is selective on the borrower's hurricane risk and absent for low-risk borrowers. A pure local-demand reallocation, in which lenders divert funds toward the disaster region to meet recovery demand (Cortés and Strahan, 2017), would predict a cut that fades with the lender's need for funds and does not sort on the unaffected borrower's own climate risk, nor the severing of high-risk relationships two years later.

The pattern is instead consistent with belief updating about climate risk. Salient events shift risk assessments even when fundamentals are unchanged (Dessaint and Matray, 2017), and Correa et al. (2026) show that after climate-linked disasters lenders reprice loans to at-risk but unaffected borrowers, a response they attribute to salience. Harvey was an extreme and highly visible event, and it arrived at the moment climate risk entered the financial

stability agenda. Google search interest in climate change around Harvey was roughly 2.5 times its level around Hurricane Sandy (Figure A6 in Appendix E). In this environment, the lenders holding the most Harvey risk had the strongest reason to reassess the hurricane exposure embedded in the rest of their portfolios, and the borrowers most affected by that reassessment were the hurricane-risky ones.

Why quantities and not prices? In the syndicated market, spreads at origination are set jointly for the syndicate, so an individual lender's reassessment of hurricane risk is more easily expressed by taking a smaller share, declining to participate, or not renewing, than by moving the price of the deal. Managing exposure through participation is also the natural instrument for a lender targeting portfolio concentration limits rather than marginal loan returns. The market-share results show the aggregate footprint of this behavior. Exposed lenders shrank precisely where the climate risk sits, and less-exposed lenders took their place. Climate risk was not eliminated from the financial system. It was reallocated toward lenders with less Harvey exposure, and the intensive and extensive margins together suggest the reallocation was deliberate and persistent rather than a transitory liquidity response.

6 Conclusion

This paper asks whether a localized climate disaster changes the supply of credit to borrowers the disaster never touched. Answering it requires separating supply from demand, and our design does so within the borrower, following Khwaja and Mian (2008). We measure each lender's exposure to Hurricane Harvey as the share of its pre-storm syndicated portfolio in the counties FEMA later declared disaster areas, and we compare lenders with different exposure lending to the same unaffected firm. The borrower fixed effect absorbs everything about the firm, including its demand for credit, so what remains is the difference in how its lenders behaved.

Exposed lenders cut credit to unaffected borrowers by 4.0 to 4.8 percent of loan volume per standard deviation of exposure when pairs are weighted by loan size, and by 2.4 to 3.4 percent when relationships count equally, with no adjustment in spreads or maturities. The contraction is unchanged when every borrower within 700 miles of the impact zone is excluded, concentrates in borrowers whose own counties carry the highest historical hurricane risk, and aggregates to a two percentage point shift of exposed lenders' portfolios out of

hurricane-risky counties, where they ceded market share to less-exposed competitors. The adjustment follows a sequence, amounts first and relationships later. By the end of 2019 exposed lenders had let high-risk relationships lapse at a rate of 1.4 percentage points per standard deviation of exposure. Placebo event dates show no contraction.

These facts select an interpretation. A lender short of capacity lends less to everyone. A lender that has revised its assessment of hurricane risk lends less to the borrowers it now considers exposed. The cut we estimate is close to zero for low-risk borrowers, appears at borrowers a thousand miles from the storm, and builds over two quarters rather than arriving with the landfall. That is the signature of a portfolio being rebalanced, not of a balance sheet running short. The timing tells us why the reassessment happened then. Harvey made landfall two months after the Financial Stability Board's Task Force on Climate-related Financial Disclosures issued the first framework asking financial institutions to measure and disclose their climate exposure (Task Force on Climate-related Financial Disclosures, 2017), and the Network for Greening the Financial System was founded that December (Network for Greening the Financial System, 2018). The storm was information about hurricane risk, and the new supervisory environment raised the cost of holding exposure that would now be visible. A lender with more Harvey-county loans both learned more from the storm and had more to account for under the emerging disclosure regime, so both forces push the same lenders in the same direction, and both are revisions of beliefs rather than losses of capacity.

Each piece of this reading has independent support. Salient events shift managers' risk assessments even when fundamentals are unchanged (Dessaint and Matray, 2017), and after climate-linked disasters lenders raise spreads on at-risk but undamaged borrowers (Correa et al., 2026), of which our quantity result is the counterpart. The magnitude matches the roughly 3 percent contraction that Ivanov et al. (2022) estimate for distant borrowers using within-county variation, and the sorting on the borrower's own disaster risk echoes Rehbein and Ongena (2022), who find that banks hit by a flood retrench most from the regions they anticipate as future disaster risks. A reallocation driven by rebuilding demand instead draws its map from the disaster region's funding needs (Cortés and Strahan, 2017). The map here is drawn by the destination's hurricane risk.

The estimates come with a boundary. The within-borrower design needs firms that borrow from several lenders at once, and the firms that identify our estimates average 6.1 of them. Firms like these can replace a retreating lender, draw on an existing facility, or issue

in the bond market, so for them the reallocation shuffles credit across lenders more than it removes credit. Our estimates measure that reallocation. They do not measure the total credit outcome for any firm, which depends on what its other lenders did, the component the design nets out by construction.

The firms outside our sample are the ones for which the same lender behavior would bite hardest. A small firm in a hurricane-prone county typically borrows from one or two local banks, has no bond market to fall back on, and holds a relationship whose value is the lender's accumulated private information about it. For such a firm a severed relationship is not a reshuffling of credit but a loss of access, which is why lender health passed into employment for small and medium firms in Chodorow-Reich (2014) while large firms were insulated. Our extensive-margin result is the one that matters most for them. Exposed lenders did not only shave amounts, they let hurricane-risky relationships lapse, and they did so at the smallest relationships first. Read against a bank-dependent borrower base, what looks like a modest portfolio tilt in the syndicated market is a withdrawal of credit from firms with nowhere else to go, accumulating in precisely the counties where climate risk is highest.

The findings carry two implications for how climate risk in credit markets is measured. First, aggregate credit in a risky region can look stable while changing hands. The reallocation moves hurricane exposure toward lenders the storm did not reach, and whether those lenders are better placed to bear it is the question a stress test should ask, rather than whether regional credit volumes hold up. Second, pricing-based evidence understates the adjustment. Spreads and maturities do not move at all in our setting, so a study reading only the price of credit would conclude that nothing happened.

The question that remains for future research is whether the same reallocation removes credit from borrowers who cannot substitute across lenders. The firms in our sample replace a retreating lender with another. Smaller firms cannot, since they depend on one or two relationships, and those relationships are often with smaller local banks whose own portfolios are concentrated in the same hurricane-risky regions, so both sides of that market are more exposed to the reallocation we document. Small commercial loans in Community Reinvestment Act data make the question answerable, and the real cost of a withdrawn relationship should be visible there.

References

Becerra-Ornelas, Alejandro Uriel, Mariela Dal Borgo, and David José Jaume, 2026, Credit guarantees and relationship lending: Evidence from natural disasters, *Working Paper* .

Berg, Gunhild, and Jan Schrader, 2012, Access to credit, natural disasters, and relationship lending, *Journal of Financial Intermediation* 21, 549–568.

Bernstein, Asaf, Matthew T. Gustafson, and Ryan Lewis, 2019, Disaster on the horizon: The price effect of sea level rise, *Journal of Financial Economics* 134, 253–272.

Billings, Stephen B., Emily A. Gallagher, and Lowell Ricketts, 2022, Let the rich be flooded: The distribution of financial aid and distress after hurricane harvey, *Journal of Financial Economics* 146, 797–819.

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

Chu, Yongqiang, Saiying Deng, and Cong Xia, 2020, Bank Geographic Diversification and Systemic Risk, *The Review of Financial Studies* 33, 4811–4838.

Collier, Benjamin, Sabrina T. Howell, and Lea Rendell, 2024, After the Storm: How Emergency Liquidity Helps Small Businesses Following Natural Disasters, Working paper, Social Science Research Network.

Conlon, Thomas, Rong Ding, Xing Huan, and Zhifang Zhang, 2024, Climate risk and financial stability: evidence from syndicated lending, *The European Journal of Finance* 30, 2001–2031.

Correa, Ricardo, Ai He, Christoph Herpfer, and Ugur Lel, 2026, The Rising Tide Lifts Some Interest Rates: Climate Change, Natural Disasters, and Loan Pricing, *Journal of Finance* .

Cortés, Kristle Romero, 2014, Rebuilding after Disaster Strikes: How Local Lenders Aid in the Recovery, Working paper, Social Science Research Network.

Cortés, Kristle Romero, and Philip E. Strahan, 2017, Tracing out capital flows: How financially integrated banks respond to natural disasters, *Journal of Financial Economics* 125, 182–199.

De Marco, Filippo, and Nicola Limodio, 2024, The financial transmission of a climate shock: El Niño and US banks, Working paper, January 31, 2024. Available at SSRN: https://ssrn.com/abstract=4279194.

Dessaint, Olivier, and Adrien Matray, 2017, Do managers overreact to salient risks? Evidence from hurricane strikes, *Journal of Financial Economics* 126, 97–121.

Froot, Kenneth A., 2001, The market for catastrophe risk: A clinical examination, *Journal of Financial Economics* 60, 529–571.

Froot, Kenneth A., David S. Scharfstein, and Jeremy C. Stein, 1993, Risk management: Coordinating corporate investment and financing policies, *The Journal of Finance* 48, 1629–1658.

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, 55–82.

Gallagher, Justin, and Daniel A. Hartley, 2025, Natural Disasters, Local Bank Market Share, and Economic Recovery, Working paper, Social Science Research Network.

Garmaise, Mark J., and Tobias J. Moskowitz, 2009, Catastrophic Risk and Credit Markets, *The Journal of Finance* 64, 657–707.

Gelman, Andrew, and David K. Park, 2009, Splitting a predictor at the upper quarter or third and the lower quarter or third, *The American Statistician* 63, 1–8.

Giglio, Stefano, Matteo Maggiori, Krishna Rao, Johannes Stroebel, and Andreas Weber, 2021, Climate change and long-run discount rates: Evidence from real estate, *The Review of Financial Studies* 34, 3527–3571.

Giglio, Stefano, Matteo Maggiori, and Johannes Stroebel, 2015, Very Long-Run Discount Rates, *The Quarterly Journal of Economics* 130, 1–53.

Goetz, Martin R., Luc Laeven, and Ross Levine, 2016, Does the geographic expansion of banks reduce risk?, *Journal of Financial Economics* 120, 346–362.

Goldsmith-Pinkham, Paul, Matthew T Gustafson, Ryan C Lewis, and Michael Schwert, 2023, Sea-Level Rise Exposure and Municipal Bond Yields, *The Review of Financial Studies* 36, 4588–4635.

Hart, Oliver D., and Dwight M. Jaffee, 1974, On the application of portfolio theory to depository financial intermediaries, *The Review of Economic Studies* 41, 129–147.

Holmström, Bengt, and Jean Tirole, 1997, Financial intermediation, loanable funds, and the real sector, *The Quarterly Journal of Economics* 112, 663–691.

Hosono, Kaoru, Daisuke Miyakawa, Taisuke Uchino, Makoto Hazama, Arito Ono, Hirofumi

Uchida, and Iichiro Uesugi, 2016, Natural disasters, damage to banks, and firm investment, *International Economic Review* 57, 1335–1370.

Ivanov, Ivan T., Marco Macchiavelli, and João A. C. Santos, 2022, Bank lending networks and the propagation of natural disasters, *Financial Management* .

Ivashina, Victoria, and David Scharfstein, 2010, Bank lending during the financial crisis of 2008, *Journal of Financial Economics* 97, 319–338.

Javadi, Siamak, and Abdullah-Al Masum, 2021, The impact of climate change on the cost of bank loans, *Journal of Corporate Finance* 69, 102019.

Jiménez, Gabriel, Steven Ongena, José-Luis Peydró, and Jesús Saurina, 2012, Credit supply and monetary policy: Identifying the bank balance-sheet channel with loan applications, *American Economic Review* 102, 2301–2326.

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

Koetter, Michael, Felix Noth, and Oliver Rehbein, 2020, Borrowers under water! Rare disasters, regional banks, and recovery lending, *Journal of Financial Intermediation* 43, 100811.

Markowitz, Harry, 1952, Portfolio selection, *The Journal of Finance* 7, 77–91.

Mestieri, Martí, Paula Bustos, Jacopo Ponticelli, and Christoph Albert, 2026, Climate change and the spatial reallocation of capital, Working paper, April 7, 2026.

Murfin, Justin, and Matthew Spiegel, 2020, Is the Risk of Sea Level Rise Capitalized in Residential Real Estate?, *The Review of Financial Studies* 33, 1217–1255.

Network for Greening the Financial System, 2018, First progress report.

NOAA, 2025, U.s. billion-dollar weather and climate disasters, Accessed: 2026-04-09.

Ouazad, Amine, and Matthew E. Kahn, 2022, Mortgage Finance and Climate Change: Securitization Dynamics in the Aftermath of Natural Disasters, *The Review of Financial Studies* 35, 3617–3665.

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* 90, 30–45.

Pyle, David H., 1971, On the theory of financial intermediation, *The Journal of Finance* 26, 737–747.

Rehbein, Oliver, and Steven Ongena, 2022, Flooded through the back door: The role of

bank capital in local shock spillovers, *Journal of Financial and Quantitative Analysis* 57, 2627–2658.

Schüwer, Ulrich, Claudia Lambert, and Felix Noth, 2019, How do banks react to catastrophic events? evidence from hurricane katrina, *Review of Finance* 23, 75–116.

Task Force on Climate-related Financial Disclosures, 2017, Recommendations of the task force on climate-related financial disclosures: Final report, Technical report, Financial Stability Board.

Appendix

This appendix collects the supporting material. Appendix A reports descriptive statistics, Appendix B derives the portfolio model's first-order condition, Appendix C shows the distribution of lender exposure to Harvey, Appendix D collects the unweighted counterparts of the main exhibits together with additional results, and Appendix E contains the robustness and interpretation material.

A Descriptive statistics

Table A1: Historical Hurricane Exposure of U.S. Counties

	Mean	25th	Median	75th	Max
Panel A. All counties (N = 3,233)					
Hurricane passages within 100 miles	19.1	1	13	28	107
Hurricane passages within 50 miles	7.3	0	4	11	50
Panel B. Counties of sample borrowers (N = 4,118 borrowers)					
Hurricane passages within 100 miles	22.9	5	18	35	103

Notes: County-level exposure computed from NOAA HURDAT2 storm tracks, 1851 to the day before Harvey's landfall, for the continental United States. A passage is a storm at hurricane intensity within the stated distance of the county centroid. Panel B repeats the first row across the 4,118 borrowers in the sample, so counties enter once per borrower.

Table A2: Lender–Borrower Relationships by Margin

	All pairs		Drop Harvey		Drop Harvey+Irma	
	N	Share	N	Share	N	Share
Both pre and post (intensive margin)	5,512	0.21	4,815	0.21	4,351	0.21
Pre only (relationship ends)	12,785	0.48	11,302	0.48	10,147	0.49
Post only (new relationship)	8,322	0.31	7,286	0.31	6,415	0.31
Total	26,619	1.00	23,403	1.00	20,913	1.00

Notes: Pairs built from the tranche-level panel before the intensive-margin restriction. A pair is a lender parent × borrower with at least one active tranche in the stated window. "Drop Harvey" and "Drop Harvey+Irma" remove pairs whose borrower is located in a Harvey (or Harvey or Irma) FEMA county.

Table A3: Sample Composition by Lender Type and Borrower Group

	Entity level (unique)		Tranche–lender level	
	Count	(%)	Count	(%)
Panel A. Lender type				
US bank	467	40.8	30,059	58.3
Foreign bank	321	28.1	14,469	28.1
Finance company	154	13.5	2,262	4.4
Investment bank	28	2.4	2,560	5.0
Other	174	15.2	2,204	4.3
Total lenders (banks)	1,144 (788)	100 (68.9)	51,554 (44,528)	100 (86.4)
Panel B. Broad borrower group				
Corporates	3,373	81.9	42,144	81.7
Non-bank financials	363	8.8	3,664	7.1
Utilities	220	5.3	3,133	6.1
Media/communications	144	3.5	2,469	4.8
Other	18	0.4	144	0.3
Panel C. Borrower industry (largest groups)				
Financial services	371	9.0	3,735	7.2
General manufacturing	324	7.9	3,872	7.5
Healthcare	274	6.7	3,552	6.9
Technology	259	6.3	4,375	8.5
Oil and gas	251	6.1	2,887	5.6
Other industries	2,639	64.1	33,133	64.3
Total borrowers	4,118	100	51,554	100

Notes: Whole sample, hurricane quarter excluded. In Panel A the entity level counts the 1,144 unique lenders; in Panels B and C it counts the 4,118 unique borrowers, and the two panels partition the same totals. Lender types come from DealScan's institution classification, with US bank including thrifts and each parent classified by the modal type of its entities. "Other" in Panel A aggregates direct lenders, investment funds, insurance companies, nonfinancial corporates, and unclassified lenders. "Other" in Panel B includes banks and government entities. "Other industries" in Panel C aggregates all industries with under six percent of borrowers. Where DealScan revises a tranche's terms the earliest record is kept, so a tranche already active enters with its pre-Harvey terms.

Table A4: Sample Composition by Loan Type and Purpose

	Tranche level (unique)		Tranche–lender level	
	Count	(%)	Count	(%)
Panel A. Loan type				
Term loan	4,562	51.8	22,999	44.6
Revolver	3,810	43.3	26,224	50.9
Other types	428	4.9	2,331	4.5
Panel B. Loan purpose				
General purpose	5,554	63.1	36,608	71.0
Acquisition	791	9.0	3,117	6.0
Leveraged buyout	681	7.7	2,363	4.6
Takeover	423	4.8	2,655	5.1
Dividend recapitalization	302	3.4	1,126	2.2
Real estate loan	221	2.5	828	1.6
Sponsored buyout	220	2.5	748	1.5
Working capital	162	1.8	1,254	2.4
Merger	60	0.7	397	0.8
Other purposes	386	4.4	2,458	4.8
Total	8,800	100	51,554	100
Panel C. Activity records (tranche–lender–quarter)				
Originations			17,761	27.1
Amendments			47,710	72.9
Total			65,471	100

Notes: Whole sample, pre- and post-years, hurricane quarter excluded. The tranche-lender level counts each lender's participation in a tranche once. Panels A and B partition the same 8,800 tranches and 51,554 tranche-lender observations. "Other types" includes bridge loans, letters of credit, and minor facilities. "Other purposes" aggregates project finance, spinoffs, recapitalizations, exit and debtor-in-possession financing, and residual categories. Panel C counts quarterly activity records before collapsing to tranche-lender pairs. Where DealScan revises a tranche's terms the earliest record is kept, so a tranche already active enters with its pre-Harvey terms.

Table A5: Descriptive Statistics. Pre- and Post-Harvey Windows

	Pre-year (Jul 2016–Jun 2017)	Post-year (Oct 2017–Sep 2018)
Panel A. Scale		
Tranches active	6,074	4,435
Total volume (USD bn)	2,358	1,855
Borrowers	3,008	2,111
Lender parents	564	455
Lender–borrower pairs	18,297	13,834
Panel B. Credit conditions (tranche level)		
Tranche amount, mean (USD mn)	388	418
Tranche amount, median (USD mn)	140	150
Margin over LIBOR, mean (bps)	334	319
Margin over LIBOR, median (bps)	300	275
Tenor, mean (months)	54.4	53.0
Tenor, median (months)	60	59
Lenders per tranche, median	4	4
Multi-lender tranches (%)	92	89
Panel C. Loan type and purpose (% of tranches)		
Term loans	52	52
Revolvers	43	44
General purpose	64	70
Acquisition-related	24	21
Panel D. Lender type (% of lender appearances)		
US banks	55	55
Foreign banks	32	32
Nonbanks	11	12
Panel E. Borrowers in Harvey FEMA counties		
Share of tranches (%)	9.9	10.5
Share of volume (%)	12.5	12.4

Notes: DealScan syndicated loan sample, whole sample before any intensive-margin restriction. Pre-year July 2016 through June 2017; post-year October 2017 through September 2018; the hurricane quarter (2017Q3) is excluded. A tranche with lending activity in both windows appears in both columns. Lenders are aggregated to the parent institution; missing syndicate shares are imputed following Chodorow-Reich (2014). Acquisition-related combines acquisition, takeover, merger, and buyout purposes. DealScan records revised terms under the same tranche identifier. Each column is deduplicated on its own, so where a tranche appears in more than one quarter of a window the earliest record in that window is kept and the post-year column reports post-storm terms.

Table A6: Pre- Versus Post-Harvey Loan Conditions by Borrower Hurricane Risk

	Hurricane-prone counties		Less-prone counties	
	Pre	Post	Pre	Post
Tranches	3,963	2,917	2,111	1,518
Borrowers	1,960	1,385	1,050	726
Total volume (USD bn)	1,587	1,241	771	614
Median tranche amount (USD mn)	140	150	135	154
Median margin (bps)	300	275	300	275
Median tenor (months)	60	60	60	58

Notes: Whole sample; windows and tranche-level deduplication as in Table A5. Hurricane-prone: the borrower's county records at least 10 hurricane passages within 100 miles between 1851 and Harvey (NOAA HURDAT2); counties with no recorded passage information are classified as less prone. Volume and medians are computed over unique tranches.

B Derivation of the lender's first-order condition

This section derives equation (3) from the problem stated in equation (2), then derives the prediction in equation (5) by differencing the two periods, and closes by mapping it to the estimating equation (13).

Portfolio moments. Start from the return in equation (1) and add up the lender's positions,

$$\sum_i L_{ij} R_i \;=\; \sum_i L_{ij} r_i \;-\; \lambda x \sum_i L_{ij} h_i \;+\; \sum_i L_{ij}\, \varepsilon_i. \tag{A1}$$

The middle term collapses to a single scalar, $H_j \equiv \sum_i L_{ij} h_i$, the lender's aggregate climate exposure. This is the step that gives the model its structure. Each borrower's climate loading reaches the portfolio only through this one number, so the lender's whole hurricane position can be summarized by it.

Take moments under the lender's own beliefs, with $E_j[x] = \theta_j$ and $\mathrm{Var}_j(x) = \sigma^2$, and

with the ε_i independent of x and of each other with common variance s^2,

$$E_j\Big[\textstyle\sum_i L_{ij}R_i\Big] = \displaystyle\sum_i L_{ij}r_i \;-\; \lambda\,\theta_j H_j, \tag{A2}$$

$$\mathrm{Var}_j\Big(\textstyle\sum_i L_{ij}R_i\Big) = \lambda^2\sigma^2 H_j^2 \;+\; s^2\displaystyle\sum_i L_{ij}^2. \tag{A3}$$

The variance splits into two pieces. The first is systematic and grows with the square of the lender's total climate exposure, so it is the part that holding many hurricane-prone loans cannot diversify away. The second is idiosyncratic and depends only on how large the individual positions are.

The optimum. Attach a multiplier $\mu_j \geq 0$ to the capacity constraint and write the Lagrangian for the problem in equation (2),

$$\mathcal{L}_j \;=\; \sum_i L_{ij}\big(r_i - \lambda\theta_j h_i\big) \;-\; \frac{\gamma}{2}\Big[\lambda^2\sigma^2 H_j^2 + s^2\sum_i L_{ij}^2\Big] \;-\; \mu_j\Big(\sum_i L_{ij} - K_j\Big). \tag{A4}$$

Now differentiate with respect to one position L_{ij}. The only subtlety is that L_{ij} enters the systematic variance through H_j, and $\partial H_j/\partial L_{ij} = h_i$, so that term contributes $\gamma\lambda^2\sigma^2 H_j h_i$ after the factor of two from the square cancels the one-half in the objective. This is why the covariance charge carries no fraction. The condition is

$$\underbrace{r_i - \lambda\theta_j h_i}_{\text{expected return}} \;-\; \underbrace{\gamma\lambda^2\sigma^2 H_j h_i}_{\text{covariance charge}} \;-\; \underbrace{\gamma s^2 L_{ij}}_{\text{own variance}} \;-\; \underbrace{\mu_j}_{\text{capacity}} \;=\; 0, \tag{A5}$$

and solving it for L_{ij} gives equation (3).

Equation (A5) characterizes the optimum and not merely a stationary point. Collect the loadings in the vector $h = (h_i)_i$. The Hessian of the objective is $-\gamma\big(s^2 I + \lambda^2\sigma^2 hh'\big)$, which is negative definite whenever $s^2 > 0$, so the objective is strictly concave and the solution is unique. The constraint is linear, so the Kuhn–Tucker conditions $\mu_j \geq 0$ and $\mu_j\big(\sum_i L_{ij} - K_j\big) = 0$ complete the characterization. We work at an interior optimum throughout and do not impose $L_{ij} \geq 0$.

Solving out the aggregate exposure. Equation (3) is implicit, because H_j is itself built from the positions it determines. It can be solved in closed form. Multiply equation (A5) by h_i, sum over borrowers, and write $S_1 \equiv \sum_i h_i$, $S_2 \equiv \sum_i h_i^2$ and $S_r \equiv \sum_i h_i r_i$. Collecting the H_j terms gives

$$H_j \;=\; \frac{S_r - \mu_j S_1 - \lambda \theta_j S_2}{\gamma\big(s^2 + \lambda^2 \sigma^2 S_2\big)}. \tag{A6}$$

The denominator is strictly positive, so the fixed point exists and is unique. Equation (A6) also makes the feedback explicit. A lender that turns more pessimistic, raising θ_j, lowers its own H_j, which lowers the covariance charge on every hurricane-prone loan and so damps the very tilt that the pessimism set off. Holding H_j fixed, as the main text does, therefore overstates the tilt that a given belief revision produces. The equilibrium feedback damps it, so the tilt we estimate is if anything an understatement of the belief channel.

From the first-order condition to equation (5). The prediction follows from writing equation (3) once for each period and subtracting. In period 0 every lender holds the common prior θ^0 and has shadow cost μ_j^0, so its position in borrower i is

$$L_{ij}^0 \;=\; \frac{r_i - \mu_j^0 - \lambda \theta^0 h_i - \gamma \lambda^2 \sigma^2 H_j h_i}{\gamma s^2}. \tag{A7}$$

In period 1 Harvey has changed the two state variables. Beliefs are now θ_j^1 and the shadow cost is μ_j^1, so

$$L_{ij}^1 \;=\; \frac{r_i - \mu_j^1 - \lambda \theta_j^1 h_i - \gamma \lambda^2 \sigma^2 H_j h_i}{\gamma s^2}. \tag{A8}$$

Subtract the first from the second. The spread r_i cancels because it did not change, and the covariance term cancels because we hold H_j fixed, the second approximation discussed in the footnote to equation (5). Only the two changes survive,

$$\Delta L_{ij} \;=\; L_{ij}^1 - L_{ij}^0 \;=\; -\,\frac{\mu_j^1 - \mu_j^0}{\gamma s^2} \;-\; \frac{\lambda\big(\theta_j^1 - \theta^0\big)}{\gamma s^2}\, h_i. \tag{A9}$$

The updating rule in equation (4) sets $\theta_j^1 - \theta^0 = \kappa(e_j)\Delta$. Substituting it and naming the two pieces

$$\phi(e_j) \;\equiv\; \frac{\mu_j^1 - \mu_j^0}{\gamma s^2}, \qquad\qquad \psi(e_j) \;\equiv\; \frac{\lambda\, \kappa(e_j)\, \Delta}{\gamma s^2}, \tag{A10}$$

gives $\Delta L_{ij} = -\phi(e_j) - \psi(e_j)\, h_i$ in dollars.

The signs of the two slopes carry the content of the model. The storm destroyed more capital at lenders with more of their book in the damaged counties, so the constraint tightened more and μ_j rose more for them, which gives $\phi' > 0$. Exposed lenders also put more weight on the news, $\kappa' > 0$, and the news was bad, $\Delta > 0$, which gives $\psi' > 0$. The same separation can be read directly from equation (3) by differentiating it holding H_j fixed,

$$\frac{\partial L_{ij}}{\partial \mu_j} \;=\; -\frac{1}{\gamma s^2}, \qquad\qquad \frac{\partial L_{ij}}{\partial \theta_j} \;=\; -\frac{\lambda h_i}{\gamma s^2}. \tag{A11}$$

A tighter capacity constraint cuts every position by the same number of dollars and makes no reference to h_i at all. A pessimistic revision cuts positions in proportion to h_i and leaves a borrower with $h_i = 0$ untouched.

Two final steps bridge the dollar difference in equation (A9) to the estimating equation. The outcome in the data is a log change, which relates to the dollar change through

$$\Delta \ln L_{ij} \;=\; \ln\!\Big(1 + \frac{\Delta L_{ij}}{L^0_{ij}}\Big) \;\approx\; \frac{\Delta L_{ij}}{L^0_{ij}}, \tag{A12}$$

the first-order Taylor expansion of $\ln(1+z)$ around $z = 0$, accurate when the change is small relative to the pre-period position. This is the first approximation discussed in the footnote to equation (5), and its substantive content is that the conversion is not innocuous for the capacity channel. The model says a capacity shock cuts every borrower by the same number of dollars, but the same dollar cut is a larger proportional cut at a smaller position, so in logs the capacity channel is flat across borrowers only up to variation in L^0_{ij}. It could leak into δ_H only if hurricane-prone borrowers held systematically smaller pre-storm positions. Table A15 shows they hold larger ones, so the leak biases δ_H toward zero. The borrower fixed effect α_i absorbs anything that moves borrower i's credit at all of its lenders at once, demand after the storm for example, about which the model is silent. The result is equation (5), $\Delta \ln L_{ij} \approx \alpha_i - \phi(e_j) - \psi(e_j)\, h_i$.

From equation (5) to the estimating equation. Equation (5) contains two unknown increasing functions and a continuous risk measure, while a regression estimates numbers multiplying variables, so two further approximations turn one into the other. First, approxi-

mate each function by its tangent line at zero exposure, again a first-order Taylor expansion,

$$\phi(e_j) \;\approx\; \phi_0 + \phi'\, e_j, \qquad\qquad \psi(e_j) \;\approx\; \psi_0 + \psi'\, e_j. \tag{A13}$$

Second, replace each borrower's h_i with the average of its risk group, which the group indicator writes in one line,

$$h_i \;\approx\; \bar{h}_L + (\bar{h}_H - \bar{h}_L)\,\text{High}_i, \tag{A14}$$

so a borrower with $\text{High}_i = 0$ carries $\bar{h}_L$ and a borrower with $\text{High}_i = 1$ carries $\bar{h}_H$. Substitute both into equation (5) and multiply out,

$$\begin{aligned}\Delta \ln L_{ij} \;\approx\; \alpha_i \;-\; \phi_0 \;-\; \psi_0\,\bar{h}_L \;-\; \psi_0\,(\bar{h}_H - \bar{h}_L)\,\text{High}_i \\ -\; \big[\phi' + \psi'\,\bar{h}_L\big]\,e_j \;-\; \psi'\,(\bar{h}_H - \bar{h}_L)\,e_j \times \text{High}_i.\end{aligned} \tag{A15}$$

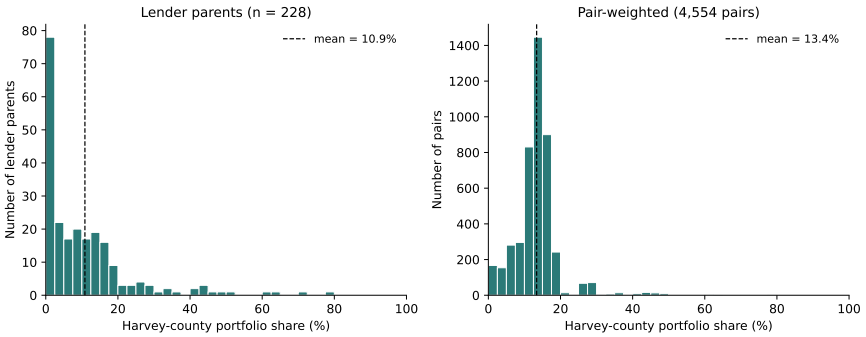
The first line varies only across borrowers. The constants are common to everyone and High_i is a fixed characteristic of borrower i, so the borrower fixed effect absorbs the entire line, which is why the estimating equation carries no separate High_i term. The second line contains exactly the two regressors of equation (13), so the regression coefficients equal the expressions multiplying them,

$$\beta \;=\; -\big[\phi' + \psi'\,\bar{h}_L\big], \qquad\qquad \delta_H \;=\; -\psi'\,(\bar{h}_H - \bar{h}_L), \tag{A16}$$

which is equation (6) in the text. Two readings follow. The capacity slope ϕ' cancels from δ_H because a capacity cut hits both risk groups identically, so the interaction responds to beliefs and only to beliefs. And β mixes the capacity slope with the belief tilt on the low-risk group's own $\bar{h}_L$, which is why belief updating predicts a β that is small and negative rather than exactly zero.

C Lender exposure to Harvey

Figure A1: Distribution of Harvey Exposure Across Lenders



Notes: Left panel: one observation per lender parent (mean 10.9%). Right panel: weighted by lender-borrower pairs in the estimation sample, the distribution the regressions use (mean 13.4%, SD 6.5 pp). Most lenders hold under 20% of their portfolio in Harvey counties, with a long right tail of locally concentrated lenders.

D Additional results

Table A7: Distance Buffers, Unweighted

Outcome: $\Delta \ln(amount)$	Pooled	>300 mi	>700 mi
Panel A: Drop Harvey borrowers (main)			
Harvey exposure (SD)	-0.034^{***}	-0.033^{**}	-0.035^{**}
	(0.013)	(0.013)	(0.014)
N	4,727	4,493	3,809
Panel B: Drop Harvey and Irma borrowers			
Harvey + Irma exposure (SD)	-0.025^{*}	-0.022^{*}	-0.026^{*}
	(0.013)	(0.013)	(0.013)
N	4,273	4,039	3,545
Borrower FE	Yes	Yes	Yes

Notes: Unweighted version of Table 4: each pair receives equal weight and the samples, buffers, and specification are identical. The buffers are cumulative exclusions, not mutually exclusive rings. In Panel B the exposure measure is the lender's pre-storm portfolio share in Harvey and Irma counties combined, standardized on the full sample. Borrower fixed effects included; standard errors clustered by lender parent in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A8: Risk Thresholds, Unweighted

	Low risk			High risk		
Outcome: $\Delta \ln(amount)$	<3	<5	<10	≥3	≥5	≥10
Harvey exposure (SD)	−0.015	−0.018	−0.013	−0.038***	−0.041***	−0.053***
	(0.030)	(0.026)	(0.019)	(0.014)	(0.012)	(0.014)
Borrower FE	Yes	Yes	Yes	Yes	Yes	Yes
N (pairs)	784	1,173	1,871	3,943	3,554	2,856
Lender parents	100	133	167	217	205	182

Notes: Unweighted version of Table 5: each pair receives equal weight and the subsamples, risk thresholds, and specification are identical. Borrower hurricane risk is the number of hurricanes passing within 100 miles of the borrower's headquarters county between 1851 and Harvey, from NOAA's HURDAT2 storm-track database. Borrower fixed effects included; standard errors clustered by lender parent in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A9: Risk Split at the Ten-Hurricane Threshold, Unweighted

Outcome: $\Delta \ln(amount)$	Low risk, <10 (base)	High risk, ≥10	Difference δ_H
Harvey exposure (SD)	−0.013	−0.053***	−0.040*
	(0.019)	(0.014)	(0.023)
N (pairs)	1,871	2,856	

Notes: Unweighted version of Table 6: each pair receives equal weight and the sample, ten-hurricane cut, and specification are identical. The low-risk group is the omitted reference; the high-risk total is computed by linear combination and coincides with the ten-hurricane column of Table A8. The difference between the two groups is −0.040 ($p = 0.09$). Borrower fixed effects included; standard errors clustered by lender parent in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A10 reports the risk-band variant of equation (13). With three groups the regression is estimated in its saturated form,

$$\Delta \ln(\text{amount})_{ij} = \alpha_i + \beta_L \left(\text{Harvey exposure}_j \times \text{Low}_i\right) + \beta_M \left(\text{Harvey exposure}_j \times \text{Middle}_i\right) + \beta_H \left(\text{Harvey exposure}_j \times \text{High}_i\right) + \varepsilon_{ij}, \tag{A17}$$

where Harvey exposure_j is standardized to unit variance and Low_i, Middle_i, and High_i mark the borrower's hurricane-risk band, counties with 0 to 6, 7 to 15, and 16 or more historical

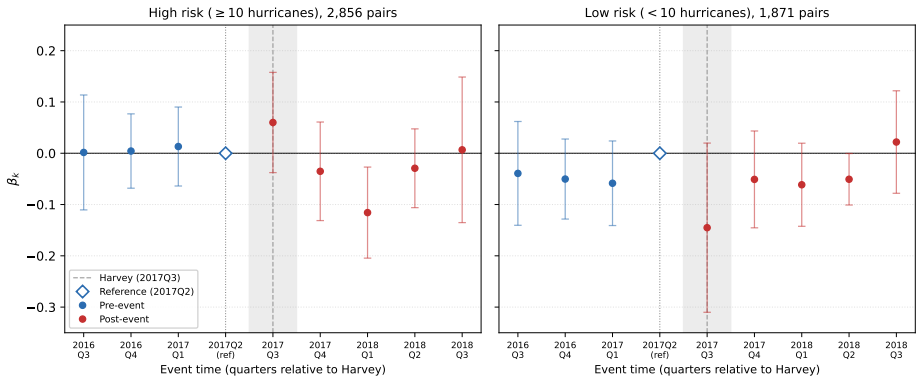
hurricanes. Because every borrower belongs to exactly one band, each coefficient is the total effect for that group rather than a difference relative to an omitted reference, and the borrower fixed effect continues to absorb demand within every group. Differences across groups are tested by linear combinations of the three coefficients. The pairwise p-values are reported in the table notes.

Table A10: Supply Contraction by Hurricane-Risk Band

Outcome: $\Delta \ln(amount)$	Low risk	Middle risk	High risk
Harvey exposure (SD), dollar weighted	−0.027	−0.061*	−0.054**
	(0.029)	(0.034)	(0.022)
N (pairs)	1,467	840	2,420

Notes: One regression of equation (A17) on all non-Harvey borrowers (4,727 pairs), grouped by the county's historical hurricane count: 0 to 6, 7 to 15, and 16 or more hurricanes within 100 miles of the centroid since 1851 (NOAA HURDAT2). Each cell is the total effect for that band. Each pair is weighted by its pre-storm loan amount as in the headline rows of Table 3. Pairwise differences across bands are not statistically significant (all $p \geq 0.41$). Borrower fixed effects included, standard errors clustered by lender parent in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

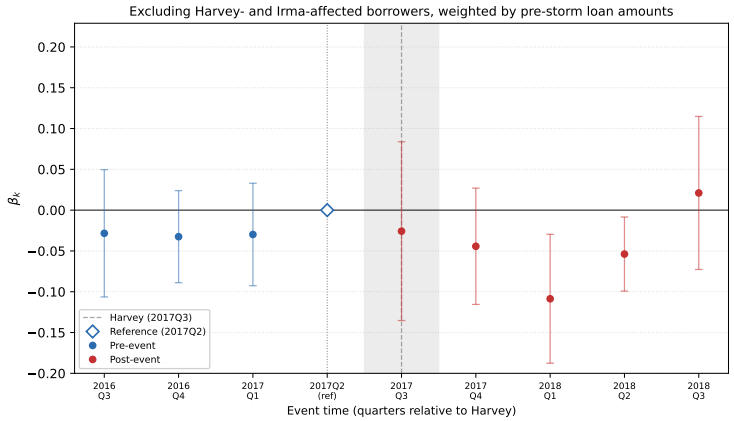
Figure A2: Event Study Split by Borrower Hurricane Risk, Dollar Weighted



Notes: Equation (11) estimated separately for borrowers above and below 10 historical hurricanes, each pair weighted by its mean quarterly pre-storm loan amount and exposure standardized before the split. High risk shows flat pre-trends (joint $p = 0.98$) and a trough of −0.116 in 2018Q1. Low risk dips noisily in the storm quarter and stays shallow afterwards. Vertical bars denote 95% confidence intervals, standard errors clustered by lender parent.

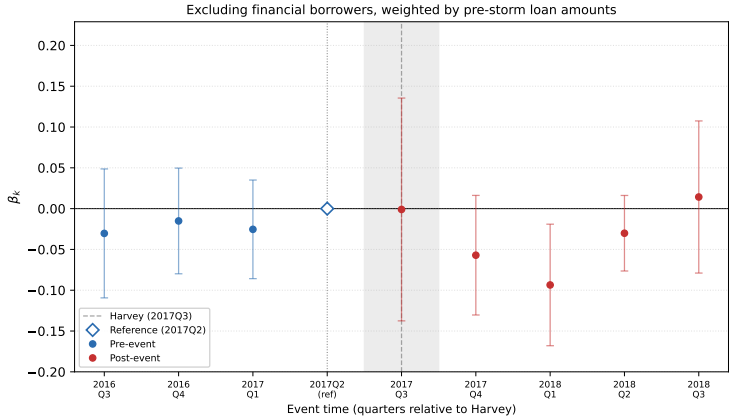
E Robustness and interpretation

Figure A3: Event Study Excluding Irma Borrowers, Dollar Weighted



Notes: Equation (11) re-estimated dropping borrowers in Harvey and Irma counties, with each pair weighted by its mean quarterly pre-storm loan amount as in Figure 3. Flat pre-trends (joint $p = 0.69$); the contraction troughs at -0.109 in 2018Q1 ($p < 0.01$). Vertical bars denote 95% confidence intervals; standard errors clustered by lender parent.

Figure A4: Event Study Excluding Financial Borrowers, Dollar Weighted



Notes: Equation (11) re-estimated dropping financial borrowers (modal SIC 6000–6999), with each pair weighted by its mean quarterly pre-storm loan amount as in Figure 3. Flat pre-trends (joint $p = 0.83$); the contraction troughs at -0.094 in 2018Q1 ($p = 0.01$). Vertical bars denote 95% confidence intervals; standard errors clustered by lender parent.

Table A11: Robustness to the Lending-Share Imputation Rule

Imputation of missing syndicate shares	$\hat{\beta}$	s.e.	p-value	Pairs
Chodorow–Reich role × structure averages (baseline)	−0.034	(0.013)	0.008	4,727
Equal split among non-reporting lenders	−0.030	(0.014)	0.037	4,727
Split by lender size in the pre-storm market	−0.025	(0.017)	0.137	4,727

Notes: Each row estimates equation (10) on the sample of Table 3, column (1): Harvey borrowers dropped, borrower fixed effects, standard errors clustered at the lender parent (242 clusters). All three rules keep the reported syndicate shares and differ only in the imputation of the missing shares. The baseline applies the Chodorow-Reich (2014) averages of reported shares by syndicate role and structure. The equal split divides the unreported residual of each tranche equally among its non-reporting lenders. The size-weighted split divides it in proportion to each non-reporting lender's number of pre-storm syndicated loan participations, a size proxy that does not use the imputed shares themselves. Lender exposure and loan amounts are both rebuilt under each rule.

Table A12: Robustness to a Recency-Weighted Measure of Hurricane Risk

Outcome: $\Delta \ln(amount)$	Lower third	Middle third	Upper third
Half-life 5 years	−0.022	−0.034*	−0.053***
	(0.020)	(0.020)	(0.019)
Half-life 10 years	−0.023	−0.032	−0.054***
	(0.020)	(0.022)	(0.019)
Half-life 20 years	−0.014	−0.034*	−0.061**
	(0.020)	(0.020)	(0.025)
Half-life 30 years	−0.014	−0.036*	−0.058**
	(0.020)	(0.021)	(0.024)
Raw count (baseline)	−0.013	−0.042**	−0.055**
	(0.021)	(0.021)	(0.023)
Pairs / lender parents		4,727 / 242	

Notes: Each row is one unweighted regression of the tercile variant of equation (13) on all non-Harvey borrowers, with terciles cut on a decayed hurricane count in which a storm y years before Harvey contributes weight $0.5^{y/h}$ instead of one. The raw count is the baseline measure of Section 2.1, equivalent to an infinite half-life. Every row cuts terciles by plain quantiles of its own measure. The middle third is the omitted reference and the lower and upper totals are computed by linear combination. Borrower fixed effects included, standard errors clustered by lender parent in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A13: Placebo Event Dates. Q3 2014, Q3 2015, and Q3 2016

Outcome: $\Delta \ln(amount)$	Pooled	>300 mi	>700 mi
Harvey 2017 (real)	-0.048^{***}	-0.047^{***}	-0.043^{**}
	(0.017)	(0.018)	(0.019)
N	4,727	4,493	3,809
Placebo 2016	0.034^{**}	0.035^{**}	0.037^{**}
	(0.015)	(0.016)	(0.018)
N	4,781	4,582	3,716
Placebo 2015	0.010	0.011	0.016
	(0.019)	(0.020)	(0.017)
N	4,282	4,104	3,447
Placebo 2014	−0.000	0.000	−0.005
	(0.015)	(0.015)	(0.018)
N	4,168	3,985	3,220
Borrower FE	Yes	Yes	Yes

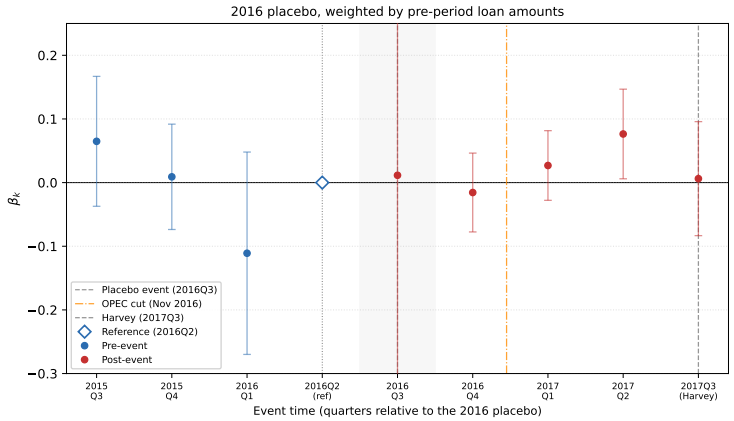
Notes: The full design of Tables 3 and 4 re-run with the event date moved to Q3 2014, Q3 2015, and Q3 2016, keeping the same Harvey counties and exposure definition. Each pair is weighted by its pre-storm loan amount as in Table 3, and the unweighted version is Table A14. The 2016 post window contains Harvey's landfall and the oil-market recovery, so it is not a clean falsification test. Borrower fixed effects included, standard errors clustered by lender parent in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A14: Placebo Event Dates, Unweighted

Outcome: $\Delta \ln(amount)$	Pooled	>300 mi	>700 mi
Harvey 2017 (real)	-0.034^{***}	-0.033^{**}	-0.035^{**}
	(0.013)	(0.013)	(0.014)
N	4,727	4,493	3,809
Placebo 2016	0.009	0.010	0.007
	(0.014)	(0.015)	(0.014)
N	4,781	4,582	3,716
Placebo 2015	0.007	0.008	0.015
	(0.012)	(0.013)	(0.011)
N	4,282	4,104	3,447
Placebo 2014	−0.009	−0.009	−0.007
	(0.010)	(0.011)	(0.013)
N	4,168	3,985	3,220
Borrower FE	Yes	Yes	Yes

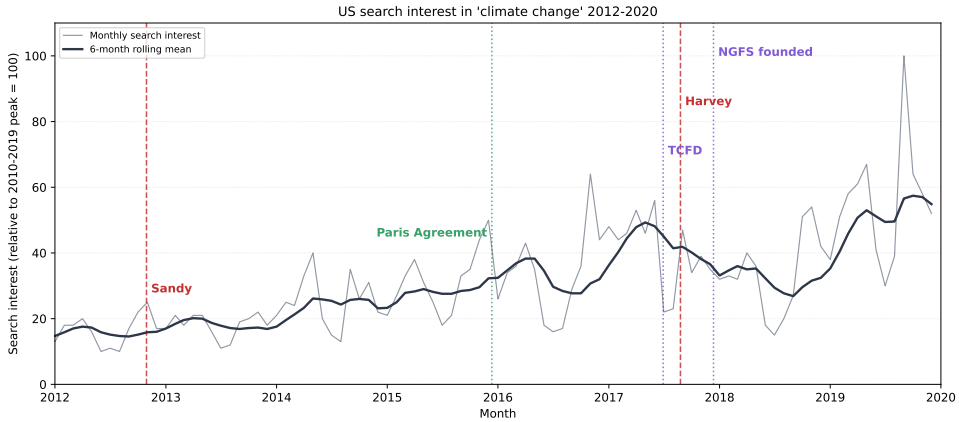
Notes: Unweighted version of Table A13, with identical samples, exposure definition, windows, and specification. Every unweighted placebo coefficient, including 2016, is small and indistinguishable from zero in every distance sample. Borrower fixed effects included, standard errors clustered by lender parent in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure A5: Dynamic 2016 Placebo. The Oil Recovery, Not a Lender Trait



Notes: Quarter-by-quarter version of the 2016 placebo in Table A13, dollar weighted, with the placebo event at 2016Q3 and 2016Q2 as the reference period. Exposure is the lender's August 2016 loan share in the future Harvey counties. No quarter before the OPEC production cut of November 30, 2016 is statistically distinguishable from zero (pre-trend joint $p = 0.21$); the path then builds to +0.075 in 2017Q2 ($p = 0.03$) as oil prices stabilized and is flat in the quarter of Harvey's landfall (+0.010). The unweighted path is flat throughout (pre-trend $p = 0.34$). Vertical bars denote 95% confidence intervals; standard errors clustered by lender parent.

Figure A6: Salience. Climate Risk Was New in 2017



Notes: Google search interest in "climate change" in the United States, 2010 to 2020. Search interest around Harvey (August 2017) was roughly 2.5 times its level around Hurricane Sandy (October 2012). The TCFD recommendations were published in June 2017 and the NGFS was founded in December 2017.

Table A15: Pre-Period Relationship Size and the Risk Tilt

Panel A. Pre-period size across risk groups

	Low risk, <10	High risk, ≥10	Difference
Pair amount, median ($m)	89.2	115.2	
Pair amount, mean ($m)	201.3	248.7	
ln(pair amount)	4.598	4.757	0.159***
			(0.036)
ln(borrower total credit)	6.295	6.331	0.036
			(0.100)
Lenders per borrower	6.07	6.05	−0.024
			(0.319)
N (pairs)	1,871	2,856	
N (borrowers)	308	472	

Panel B. The risk tilt controlling for relationship size

	Dollar weighted		Unweighted	
Outcome: $\Delta \ln(amount)$	(1)	(2)	(3)	(4)
Harvey exposure (SD)	−0.011	0.007	−0.013	0.015
	(0.023)	(0.022)	(0.019)	(0.027)
× High risk (δ_H)	−0.064**	−0.058**	−0.040*	−0.049**
	(0.030)	(0.029)	(0.023)	(0.025)
× ln(pair amount)		−0.030**		0.036
		(0.014)		(0.028)
High-risk total	−0.074***	−0.050**	−0.053***	−0.034*
	(0.022)	(0.019)	(0.014)	(0.019)
N (pairs)	4,727	4,727	4,727	4,727

Notes: Sample and specification of Table 6, in other words all non-Harvey borrowers with borrowers split at ten historical hurricanes in their county. Panel A compares pre-period size across the two risk groups. Pair amount is the lender's pre-storm loan amount to that borrower, borrower total credit sums it across the borrower's lenders, and the difference column reports the gap in the stated units with the standard error of a two-sample t test below. Panel B adds to equation (13) the interaction of standardized Harvey exposure with the log of the pair's pre-period amount, centered on the estimation-sample mean so that the exposure coefficient is evaluated at the average relationship. Columns (1) and (3) reproduce Tables 6 and A9. The high-risk total is computed by linear combination. Borrower fixed effects included; standard errors clustered by lender parent in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Natalia Madrid

PhD Candidate in Economics, University of California San Diego
San Diego, California · namadrid@ucsd.edu · nataliamadridb.github.io

Education

Ph.D. in Economics, University of California San Diego, USA 2023–Present

Master of Arts in Economics, University of California San Diego, USA 2025

Master in Financial Management, Universidad de Concepción[1], Chile 2019

B.Sc. in Industrial Engineering, Universidad de Concepción, Chile 2018

Complementary Education

Summer School in Dynamic Structural Econometrics, University of Michigan 2026

Summer School in Environmental Economics, University of California Berkeley 2025

Research

Research Interests: Climate Finance, Banking, and Insurance

Conference and Seminar Participation

2027: Scripps–Climate Economics Workshop, UCSD (co-organizer, forthcoming)

2026: NBER Climate and Nature Finance Conference (poster); Finance Brown Bag Seminar, UCSD (presentation); Climate Conference, UCSD (presentation)

2025: Internal Macroeconomics and Environmental Economics Seminars (presentation), UCSD

2021: Chilean Network of Innovation Researchers Conference, Chile (presentation)

Published Articles

Mardones, C. and Madrid Becerra, N. (2020). "Ex-post evaluation of the R&D tax incentive law in Chile." *Academia Revista Latinoamericana de Administración*, 33(3/4), 337–359. doi.org/10.1108/ARLA-03-2019-0092

Research Experience

Summer Research Intern, Banco de México, México 2026
Summer Research Fellow at the General Directorate of Financial Stability.

Summer Graduate Research Fellow, University of California San Diego, USA 2024
Independent summer research fellowship. Faculty advisor: Prof. Judson Boomhower.

Research Assistant, Universidad de Concepción, Chile 2018
Supported research under Profs. Cristian Mardones and Pablo Catalán.

Teaching Experience

Teaching Assistant, University of California San Diego 2024–Present

- Graduate: Finance & Valuation in Corporate Finance (Prof. William Mullins) Winter 2025–26
- Graduate: Real Estate Finance (Prof. Carlos Avenancio-Leon) Spring 2026
- Undergraduate: Econometrics (Prof. Graham Elliot, Prof. Yinlin Dai) Fall 2024–25

[1]The Universidad de Concepción ranks as the 3rd best university in Chile according to the Academic Ranking of World Universities (ARWU) and is ranked 11th in Latin America by the QS World University Rankings. I graduated with the highest distinction in both undergraduate and graduate studies, ranking in the top 5% of the class.

Lecturer, Universidad de Concepción, Chile 2019–2020
Instructor for 120 students; designed course materials and supervised teaching assistants.

- Undergraduate: Accounting and Industrial Costs — First Semester
- Undergraduate: Economic Engineering — Second Semester

Relevant Professional Experience

Quality Management Specialist, Universidad de Concepción, Chile 2020–2022
Redesigned quality management systems for training programs using Power BI.

Consulting Analyst, CameronPartners Innovation Consultants, Chile 2019–2020
Provided financial management advisory to 10+ SMEs; contributed to an acceleration program for 20 SMEs under CORFO's Innovation Fund for Competitiveness (FIC).

Honors, Scholarships, and Grants

Richard Libby Research Award ($10,000), Rady School of Management, UCSD 2026

Travel Grant ($800), Economics Department, UC San Diego 2026

Travel Grant ($500), Alfred P. Sloan Foundation, UC Berkeley 2025

Summer Research Scholarship ($3,500), University of California San Diego 2024

Regents Fellowship ($27,000), University of California San Diego 2023

Fulbright Doctorate Scholarship ($100,000), Fulbright Commission–ANID, Chile 2020

Outstanding Woman Engineer Award, Universidad de Concepción, Chile 2019

Merit Scholarship for a Master degree ($2,500), Universidad de Concepción, Chile 2016

Bicentennial Scholarship for Bachelor degree ($12,000), Ministry of Education, Chile 2012

Community Service and Leadership

Fulbright Board Member, San Diego Fulbright Chapter 2026–Present

Summer Reading Group Coordinator, Climate Change Economics, UCSD 2025

Educational Mentor, Powerful Women Program, WoomUp, Chile 2020

Student Council Member, Industrial Engineering Student Association, UDEC 2016

Professional Memberships

Econometric Society · Association of Environmental and Resource Economists
Fulbright Association · American Finance Association

Languages and Skills

Languages: English (full professional proficiency), Spanish (native)
Computing: Python, STATA, C++, MySQL, LaTeX, HTML/CSS, Microsoft Office

References

Judson Boomhower jboomhower@ucsd.edu *(main advisor)*
Mark Jacobsen m3jacobsen@ucsd.edu *(committee member)*