Bank Competition and Strategic Adaptation to Climate Change

Dasol Kim* Luke M. Olson* Toan Phan†

April 27, 2026

Preliminary and incomplete. Please do not circulate.

Abstract

How do strategic interactions among banks shape adaptation to emerging climate risks? Using confidential supervisory data linked to granular flood-risk measures, we study HELOC lending after Hurricane Harvey. Banks exposed to the hurricane learn about flood risk and reduce originations in high-risk areas—but adaptation depends critically on whether competitors also adapt, consistent with strategic complementarities. When peer-exposure controls enter the specification, the individual adaptation coefficient attenuates by roughly 1.3 log points, and we validate the peer channel through permutation tests, a sign-reversing non-local placebo, and replication across 1,118 banks. These findings imply that climate stress tests evaluating banks in isolation will understate the systemic adaptation response.

*Office of Financial Research, 717 14th Street NW, Washington DC 20005.

†Federal Reserve Bank of Richmond, 701 E Byrd St, Richmond, VA 23219.

We thank Neth Karunamuni, Vy Nguyen, and Priya Sankar for their excellent research assistance. We thank Ralf Meisenzahl, Peter Tufano, Teng Wang, and Mayank Kumar for valuable discussions. We also thank participants in seminars at the Office of Financial Research, the Federal Reserve Bank of Richmond, the Federal Reserve Bank of Philadelphia, FDIC, and IMF, as well as participants in the IBEFA Conference, SFS Cavalcade North America Conference, NBER Summer Institute, SITE Conference, and the FDIC Banking Research Conference. The views and position of this paper are of the authors and do not necessarily represent those of the Office of Financial Research, the U.S. Department of the Treasury, the Federal Reserve Bank of Richmond, or the Federal Reserve System. All errors are ours alone.

JEL codes: D14, E6, G21, Q54. **Keywords:** Banks, loans, climate risk, adaptation, real estate, competition, coordination.

1 Introduction

Financial institutions are increasingly confronted with emergent risks, ranging from technological disruptions to the evolving hazards of a changing climate. These risks are inherently nonstationary: their statistical properties evolve over time, making it difficult to learn from past data and adapt effectively (Weitzman, 2009; Kozlowski et al., 2020). Unlike traditional risks, which are well-documented and incorporated into regulatory frameworks, emergent risks introduce uncertainty that challenges standard risk management models. Banks must therefore navigate the dual challenge of learning about new risks while simultaneously adapting their behavior, often in the absence of clear regulatory guidance.

This paper studies how strategic interactions among banks shape adaptation to emergent risks. When regulatory guidance is thin, risk management is not an isolated decision. A bank that considers pulling back from a risky market must also consider what its competitors will do: if rivals maintain their lending, the bank that adapts cedes market share; if rivals also pull back, the cost of adaptation is lower, and the combined pullback can move the market. This creates a strategic complementarity in adaptation: each bank's optimal response is increasing in its competitors' responses, and opens the door to equilibria in which risk management is sustained by mutual expectation rather than individual conviction. The complementarity need not be welfare-improving. During the subprime crisis, competitive pressures contributed to risk accumulation in precisely those segments where supervisory guidance was thin: securitization had limited historical precedent and, in turn, weak regulatory oversight, creating opportunities for regulatory leakage. In its post-mortem of the crisis, Financial Crisis Inquiry Commission (2011) identified competition as a key factor in the rapid expansion of these markets. That episode illustrates what is at stake when strate-

gic interactions govern risk-taking in the absence of oversight; our analysis documents how the same mechanism operates in the adaptation to emergent risks.

The underlying economics admit two offsetting strategic forces, and these forces are the empirical counterparts of the two complementarity channels we formalize in Section 8. On the one hand, some institutions may become informed about emergent risks before others, for example because exposure to a realized event provides access to data that would otherwise be costly to acquire. If informed banks curtail lending in risky markets, uninformed banks may adversely select into the riskier loans, a situation analogous to the "winner's curse." This force connects the private information that an informed bank acts on to the behavior of its peers: as more peers pull back, the mass of rejected borrowers reapplying to other lenders grows, and with it the marginal benefit of acting on one's own signal. On the other hand, an informed bank may hesitate to adapt if it expects competitors to continue lending, since doing so risks ceding business to rivals (Di Maggio et al., 2019). This second force does not depend on the bank's private information. It operates at the level of adaptation itself: when peers adapt, the cost of adapting falls, whether through reputational and supervisory pressure on banks whose behavior drifts from peer norms or through the simple reduction in market share at stake when rivals are also pulling back. Which force dominates is an empirical question, and the answer has direct consequences for how risk diffuses or concentrates within the banking sector. Our analysis shows that both forces are present in the data: banks adapt less when local competitors are less likely to adapt, and informed banks adapt more when local competitors are also exposed to the same shock, consistent with strategic complementarities sustaining a coordinated pullback. Even when dislocations of this kind arise only initially, they may have destabilizing effects, as the crisis experience suggests.

Our analysis uses the evolving distribution of natural disaster risks as a laboratory for studying strategic interactions in the adaptation to emergent risks. A key challenge is that risk beliefs are generally difficult to observe. Interpreting bank choices from outcomes alone is incomplete, as some outcomes are observationally consistent with both deliberate strategy

and a lack of awareness. If an econometrician observes a bank continuing to lend in a riskier area but lacks visibility into the bank's internal risk models, it remains ambiguous whether this reflects ignorance or strategic considerations. We address this issue using confidential bank supervisory data that allows us to look "under the hood" and directly identify individual bank beliefs from their internal risk models, tracking how those beliefs change in the face of emergent risk shocks. We join the supervisory data with high-resolution measures from the First Street Foundation that pinpoint property-level future flood risk, including in areas not directly affected by any past disaster. Our analysis focuses on bank portfolio exposures to Hurricane Harvey, which is not only one of the costliest natural disasters in U.S. history (NOAA, 2018) but was also notable for the uncertainty it created within the scientific community about the validity of existing loss models. Because the supervisory sample is limited to the eight largest bank holding companies, we complement it with an independent panel of 1,118 bank holding companies constructed from Equifax credit bureau data. The two samples are designed to be complementary: the supervisory data provide precise identification through direct observation of internal risk models, and the credit bureau data provide the cross-bank variation needed to identify and discipline the coordination channel.

We present three sets of results. First, banks exposed to Hurricane Harvey update their internal risk models to incorporate flood risk (Section 5). The updating is concentrated among exposed banks, begins shortly after the event, and does not revert over the two years following the storm. This pattern is consistent with exposed banks acquiring durable information about flood losses rather than responding to a transient salience shock.

Second, this learning translates into adaptation (Section 6). The response takes two forms. On existing loans, informed banks reduce credit limits, curtail drawdowns, raise pricing, and are more likely to take line management actions such as freezing or decreasing credit lines in high-risk areas, all consistent with intensive-margin adjustments to already-booked exposures. On new lending, informed banks curtail originations in high-risk areas, and this response operates primarily through the extensive margin: banks ceasing to originate

in these areas entirely rather than merely lending less within them. The origination response is concave in bank exposure, consistent with a learning mechanism in which the informational value of disaster exposure exhibits diminishing returns. A placebo test using a six-month pre-Harvey window confirms that these patterns do not predate the event, and the results are robust to alternative risk cutoffs, to Poisson pseudo-maximum likelihood estimation, and to excluding banks one at a time.

Third, adaptation depends on competitors (Section 7). We document two distinct competitive channels. Market concentration dampens the adaptation response: the reduction in new originations weakens in counties where the top lenders hold a larger share of the market, reflecting a higher cost of ceding market position when adaptation is unilateral. More importantly, banks are less likely to adapt when their local competitors are also less likely to do so, consistent with a strategic complementarity in adaptation.

The peer-exposure result is central to the paper, and its interpretation as a strategic interaction hinges on distinguishing it from more mundane alternatives. The most natural concern is that peer exposure simply proxies for a shared local shock: banks lending in the same county mechanically share similar Harvey exposures through their overlapping portfolios, so a negative coefficient might reflect demand shifts rather than coordination. Our evidence is inconsistent with this reading. Peer exposure constructed from the same banks' activities in *other* counties carries the opposite sign, a sharp sign reversal that a correlated-shock explanation cannot deliver. Controlling for county-level aggregate exposure strengthens rather than weakens the estimate, indicating that what matters is the *identity* of local competitors rather than the *level* of aggregate local exposure. A permutation test that randomly reshuffles bank identities within county by risk cells recovers the actual peer coefficient at the tail of the permutation distribution, ruling out market-level sorting. A capacity-reallocation test, which asks whether the peer effect is a mechanical consequence of competitors shifting lending across their other markets, leaves the peer coefficient intact while the capacity measure itself is indistinguishable from zero. And the peer effect sur-

vives controls for nonlinearities in own exposure, addressing the concern that peer exposure proxies for a misspecified own-exposure channel. Taken together, these tests are consistent only with a local, strategic interaction among competitors facing a common shock. The result replicates in credit bureau data covering over a thousand bank holding companies, a sample two orders of magnitude larger than the supervisory panel, where the cross-section is sufficient to discipline inference and to rule out that the peer channel is an artifact of the eight-bank supervisory setting.

In Section 8, we organize these findings in a parsimonious quantitative framework. Each bank chooses an adaptation intensity in a static game in which two distinct forms of strategic complementarity are at work. The first is a *level* complementarity: when peers adapt, the cost of adapting falls for every bank in the market, regardless of what that bank privately knows about flood risk. This channel has a natural microfoundation in the reputational and supervisory pressure that banks face when their behavior drifts from peer norms. The second is a *signal-dependent* complementarity: when peers adapt, the marginal benefit of acting on one's own private information rises, because the borrowers rejected by adapting peers reapply to remaining lenders, and the mass of such reapplications scales with the underlying bad-borrower share that only the bank's private signal identifies. This channel has a natural microfoundation in the adverse-selection spillover that arises from reapplication within a local market. The two complementarities arise from non-overlapping mechanisms and have distinct empirical signatures in our specifications: the level complementarity loads on the peer-exposure interaction that is shared across all banks in a market, while the signal-dependent complementarity loads on the interaction of peer exposure with the bank's own exposure. Under a contraction condition on the cost of adaptation, the game admits a unique equilibrium, which Hurricane Harvey *shifts* rather than *selects* between multiple candidates: the hurricane raises the private signals of exposed banks, and the complementarities amplify the resulting pullback through the coordination channel.

The framework generates four predictions beyond the baseline adaptation result, each

of which we verify. The adaptation response should be concave in own exposure, which we confirm in a quadratic specification. The triple interaction of high-risk, own exposure, and peer exposure should carry a negative sign, which we confirm in both the supervisory and the credit bureau samples. Peer exposure constructed from non-local markets should carry the opposite sign to the local peer interaction, reflecting the spatial locality of the adverse-selection mechanism; the observed sign reversal is inconsistent with correlated-shock explanations. And the own-exposure coefficient should attenuate substantially when the peer channel is included, reflecting equilibrium feedback from the coordination channel into the short regression that omits it; the observed attenuation is approximately one log point. Together, these predictions connect the reduced-form coefficients in our tables to the underlying structural parameters of the model, and they distinguish the coordination framework from alternative explanations that would predict different cross-restrictions. We read the combined evidence as indicating that competitive dynamics do not displace individual learning but operate alongside it: the individual adaptation incentive remains statistically significant in the combined specification estimated on the credit bureau sample, while the coordination channel accounts for a quantitatively important additional share of the equilibrium response.

By capturing both the learning process and the lending outcomes it induces, our approach provides a window into how emergent risks diffuse or concentrate within the banking sector. While exposed banks do learn and limit their exposure, strategic complementarities can amplify the scope for coordination failure in risk management. These findings have policy implications. Much of the regulatory attention to natural disaster risk has focused on the behavior of individual institutions, which is relevant for microprudential policy. Our evidence suggests that evaluating banks in isolation may understate the coordination channel that shapes adaptation in equilibrium, when banks anticipate that competitors facing similar exposures will also adjust. By making information about peer exposures and adaptation more widely available, climate risk disclosure could complement existing regulatory mandates by making the coordination more predictable, without requiring a catastrophic event to serve

as a coordination device. The welfare implications are nevertheless ambiguous: coordinated adaptation may improve bank stability while also restricting credit access for borrowers in high-risk areas, to the extent that the dynamics documented here generalize beyond the HELOC market.

Related Literature

Our paper connects two strands of research: the growing literature on natural disaster risks and financial markets, and the literature on bank competition and risk-taking.[1]

We contribute to the literature on how evolving disaster risks affect financial markets, and are the first, to our knowledge, to examine the role of strategic interactions among banks in this context. Our paper enriches the research on flooding and household finance (Bernstein et al., 2019; Baldauf et al., 2020; Murfin and Spiegel, 2020; Bakkensen and Barrage, 2022; Keys and Mulder, 2022; Bakkensen et al., 2025). Several papers use regulatory data to study disaster risks and bank lending. Meisenzahl (2023) finds that large banks reallocate portfolios away from disaster-prone counties; we document a similar pattern for HELOCs but leverage loan-level data with high-dimensional fixed effects following Khwaja and Mian (2008) to account for time-varying local demand and supply shocks. Correa et al. (2022) show that commercial loan spreads spike after disasters but revert quickly, consistent with the learning literature (Kousky, 2010; Gallagher and Hartley, 2017; Dessaint and Matray, 2017; Alok et al., 2020; Ouazad, 2022); we document longer-lived effects by directly observing updates to internal risk models. In documenting the strategic dimension of adaptation, our paper is related to Ouazad and Kahn (2021) and Blickle et al. (2024), who study mortgage lenders' responses to flood risk. Blickle et al. (2024) study first-lien originations; our analysis differs in focusing on HELOCs (which lack a liquid secondary market, so banks retain flood risk exposures), in directly observing internal risk model updates rather than inferring beliefs from behavior, and in documenting the role of strategic complementarities

[1] Our findings are also relevant to the literature on investment under uncertainty and to spatial economics; see Section 8 for discussion.

among competitors. A broader limitation of studies that rely on behavioral outcomes alone is the difficulty of distinguishing deliberate adaptation from other responses such as portfolio rebalancing or regulatory pressure; our access to bank risk models allows us to identify the learning channel directly. In studying how heterogeneous beliefs about climate change affect financial markets, our paper is complementary to Bakkensen et al. (2025), who analyze the asset pricing implications of disagreement about climate risk.

We also contribute to the literature on how competition affects bank risk-taking (Keys et al., 2009; Drechsler et al., 2017; Egan et al., 2017; Whited et al., 2021; Wang et al., 2022). Our analysis differs in three ways. First, we directly observe bank risk beliefs and study how they evolve over time. This allows us to distinguish actions undertaken by banks to deliberately take on risk from decisions made in the absence of awareness. Second, we study a setting with limited regulatory oversight, where the scope for strategic behavior is greater. Third, we document strategic complementarities in adaptation and provide a structural framework that distinguishes two channels, a level complementarity and a signal-dependent complementarity, connecting our evidence to the theoretical literature on coordination and financial stability (Bryant, 1980; Diamond and Dybvig, 1983; Morris and Shin, 1998; Rochet and Vives, 2004; Goldstein and Pauzner, 2005; Vives, 2014). In analyzing adaptation spillovers, our paper is also related to Di Maggio et al. (2019), who document a race to the bottom in complex mortgage issuance. We study a different margin, namely peer exposures to a common shock within local markets, and document complementarity, rather than competition, in risk-taking.

2 Institutional Background

This section provides an overview of the institutional context relevant to our analysis. The Safety and Soundness framework (12 CFR Part 30, Appendix A) serves as the foundation for supervisory oversight (Board of Governors, SR 96-14, 1996). Bank examiners assess emerging

risks within the scope of existing regulations, applying established supervisory principles (SR 96-14, 1996; OCC, Bank Supervision Process, 2019).

However, when risks do not align directly with existing frameworks, examiners may have limited authority to issue formal corrective actions without additional regulatory guidance (12 CFR Part 30, Appendix A). While the supervisory framework is designed to be adaptable, it may not explicitly address all emerging risks, as acknowledged by regulators (Federal Reserve, SR 11-7, 2011). In cases where historical data is limited, risk assessment can be particularly challenging. For instance, financial shocks stemming from novel market developments—such as the subprime mortgage-backed securities (MBS) market in the mid-2000s—have demonstrated the difficulties of evaluating risks without sufficient historical precedent (IOSCO, 2008). Another feature of the framework is to allow regulators to use best practices as a standard for risk management practices. Again, the lack of historical data associated with most emergent risks hinders the use of best standards in supervisory actions, as it inhibits the establishment of best practices amongst banks.

While these issues are generally applicable for most emergent risks, they have been particularly relevant for disaster-related risks. Supervisors have long attempted to address disaster-related risks within existing regulatory frameworks. However, these efforts have been primarily limited to guidance on disaster preparedness and post-disaster recovery. Typically, the Federal Reserve Board of Governors, the Federal Deposit Insurance Corporation, and the Office of the Comptroller of the Currency issue joint guidance to banks operating in the affected areas. Even still, the guidance is generally focused on customer access to banking services and borrower accommodations, and not necessarily on risk management practices. Moreover, the narrow scope of oversight on these matters to areas that have recently experienced a major disaster or areas that are prone to them leaves supervisors with significantly less discretion elsewhere. This is due in part to the fact that there is no recent precedent of a U.S. bank failure directly attributable to a natural disaster.

More recently, federal regulators have introduced principles for integrating disaster-

related financial risks into governance, risk management, and scenario analysis (OCC, Federal Reserve, FDIC, 2023). These principles outline expectations for large financial institutions while broader efforts to develop standardized regulatory approaches remain in the exploratory stage (Brunetti et al., 2024). However, as of 2024, which is after the sample period associated with the analysis, formal guidance has yet to be enacted.

3 Data

3.1 Data Sources

We use confidential bank supervisory data obtained from Schedule B.1 of the Federal Reserve Board's FR Y-14M for the period between January 2014 and December 2019. The regulators collect the data for the purpose of stress testing the largest U.S. bank holding companies. The Y-14M data are reported monthly by large banks who hold or service domestic home equity loans and home equity lines of credit (HELOC), and include loan-level information on each of the banks' HELOC portfolio.

One reason why our analysis focuses on the HELOC market is that it provides an excellent setting to study adaptation to emergent risks due to several institutional details. Like first-lien mortgage loans, HELOCs are collateralized by real estate assets, which are particularly exposed to future natural disaster risks due to their durability and immobility. However, unlike mortgages, HELOCs cannot be readily securitized and sold off the banks' balance sheets due to a relatively limited secondary market. HELOCs represent long-term investments that sit on bank balance sheets for quite some time, making it more likely for banks to be cognizant of associated risks. Hence, focusing on HELOCs helps us better identify the ways that banks may adapt to natural disaster risks. The focus on HELOCs also alleviates potential selection issues associated with the option to originate and sell a loan to a secondary market where market participants tend to have heterogeneous information and beliefs (Ouazad and Kahn, 2021; Bakkensen et al., 2025).

The Y-14M dataset features almost 180 million observations of 6.9 million HELOC borrowers located across 2,962 counties. To be included in the analysis, we require the availability of certain fields, including those related to a bank's internal risk model, namely each loan's default probability, and other important information including the loan terms, loan utilization, and borrower characteristics at origination. The dataset also provides each loan a unique identifier, allowing us to track the loan over time. Importantly, the data include the exact address of the underlying collateral and hence allows us to exactly measure the collateral's exposure to future flood risk, which is known to vary very granularly across space.

To assess the representativeness of the data to the complete market, we compare the loan coverage in the FR Y-14M data to a comprehensive sample based upon another data source provided by the Credit Bureau. Figure 1a shows the geographic distribution of home equity line balances represented in the bank regulatory data, by county. The choropleth is formatted such that lower (higher) densities are shaded blue (orange). The figure indicates broad geographic coverage of the dataset. We next obtain data from Equifax that provides home equity loan stock by county issued by any credit institution. The information is displayed in Figure 1b. It is formatted similarly to Figure 1a. The loan distributions between our sample and the Equifax sample are highly correlated, or 92.8%, confirming the representativeness of the confidential data. This makes sense as the reporting banks in our sample account for a majority share of approximately 59% of the entire HELOC market.

For the event study, we focus on Hurricane Harvey. Hurricane Harvey has been estimated to have caused approximately $125 billion in damages, making it one of the costliest natural disasters in U.S. history (NOAA, 2018). An important reason why the analysis focuses on Hurricane Harvey is the extent of indicators that signaled structural shifts in the distribution of natural disaster risks (Risser and Wehner, 2017). This uniquely indicated the intensification of disasters (Trenberth et al., 2018) and damage from inland flooding (Zhang et al., 2018). In contrast, coastal flooding generally accounted for a bulk of damages for previous disasters. This spurred discussion in the scientific community on the appropriateness of ex-

isting models to project loss distributions due to disasters going forward. These discussions were also relevant to financial institutions, such as banks and insurers, as many of them relied on the same models.

We obtain detailed satellite imageries that provide information on the path of Hurricane Harvey from the National Oceanic and Atmospheric Administration (NOAA) HURDAT2, using the Atlantic Hurricane Database. The data provides detailed information on the precise coordinates, maximum winds, central pressure, and the size of each tropical and subtropical cyclones in six-hour intervals. We collect the data from August 16, 2017, through September 9, 2017, to identify the ZIP codes that intersect with Hurricane Harvey's path based upon the wind radii of its focal point.

Figure 2 displays a visualization of areas affected by Hurricane Harvey. Similar to Ouazad and Kahn (2021), we focus on portions of the hurricane path that exhibit tropical storm intensities or stronger, or at least 34 knot wind speeds. This excludes the parts of Louisiana that were on the hurricane path. We impose this restriction in order to capture the most powerful portion of the storm associated with the bulk of the damages. We use the data to identify properties associated with loans in the bank regulatory data that were crossed by the hurricane path.

We also merge the Y-14M data with data from the First Street Foundation (FSF) in order to pinpoint each property's future flood risks. The FSF provides property-level projections of future flood risk, called the flood factor. The factor is defined on a scale from 1 to 10, where houses with a factor of 1 have minimal flood risk while those with a factor of 10 have very substantial risk. The flood factor is a composite score reflecting both the severity and cumulative likelihood of flooding over a 30-year period from 2021 to 2050. The factor is generated by FSF's flood model, which considers four major flood contributors: rainfall, river overflow, high tide, and coastal storm surge. The model takes into account hyper-local geographic characteristics (such as elevation, slope, and ground surface perviousness) and existing community flood protection measures (such as dunes, wetlands, and seawalls).

Importantly, the model is forward-looking, explicitly considering projected effects related to the evolving distribution of natural disaster risks, including sea-level rise and changes in precipitation. FSF also provides the proportion of properties in a 5-digit ZIP code that are classified by the Federal Emergency Management Agency (FEMA) as being in a Special Flood Hazard Area (SFHA).

3.2 Data Description

Table 1 displays the summary statistics. The loan-level probability of default measure is highly right skewed; it has a mean of 4.5% and a median of 0.2%. In the analysis, we address this issue by using the natural log transformation of one plus the default probability in points. Additionally, the FSF flood factor also exhibits right skewness and is transformed similarly for the analysis.

Figure 3 displays the geographic distribution of the flood risk measure by ZIP code. The choropleth is formatted such that lower (higher) risk levels are shaded blue (red).

Table 2 displays the histogram of the flood factor. A vast majority of the properties (85.7%) are rated as the lowest risk level. In contrast, those with score above five account for 5.9% of all properties. The table also displays the average proportion of homes within a ZIP code that are located in an official SFHA. As one may expect, this fraction increases monotonically in the underlying flood factor. The highest flood factor is associated with ZIP codes where an average of 34.4% of homes are located in an SFHA. The correlation between the flood factor and SFHA measure is 63.4%. The table also displays the coefficient of variation of the SFHA measure for each risk score. For all values, the coefficient of variation exceeds 1.0, and is pronounced for lower flood factor values.

There is considerable heterogeneity in the bank-level exposures to Hurricane Harvey. More importantly for the analysis, there is geographic variation within and across markets for which banks with and without exposures to Hurricane Harvey operate. Figure 4 displays counties where all banks have no exposure (light grey), there are both banks with and

without exposures (dark grey), and all banks with exposure (orange). Texas is excluded in our analysis for reasons discussed in the next section, and hence Texas is not included in the chart. Most counties feature both banks with and without exposures. Most of the counties where all banks have exposures are concentrated in midwestern and northeastern states.

4 Empirical Design

4.1 Learning and Exposures to Natural Disaster

An important identifying assumption underlying our main tests on competition relates to how banks learn about future flood risks following the realization of a natural disaster. Specifically, we conjecture that banks with larger ex ante exposures to Hurricane Harvey are more likely to learn and therefore are better able to incorporate future flood risk into their internal risk models. If so, we will be able to directly trace changes in a bank's beliefs about natural disaster risks based upon the bank's exposures to the hurricane. Furthermore, tracking bank beliefs will allow us to better identify adaptation responses and how they are conditioned by competitive factors.

Intuitively, some banks may have better access to information about risks that would be otherwise costly to obtain relative to other banks due to exposures to natural disasters. Following a natural disaster, banks may be able to collect data on affected properties associated with their loan portfolios. On the one hand, the occurrence of a natural disaster may inform the frequency of such disasters from happening again. This type of information is highly localized and so applicable only to borrowers in specific regions. On the other hand, the severity of the damages due to the disaster may also provide insights into losses associated with a broader set of events. For example, the damage during a hurricane may provide information about losses due to heavy rainfalls in other regions as well. The magnitude of the losses, in turn, may directly inform the borrower's ability to repay in the event of such events, even in areas unaffected by the disaster. This form of soft information may provide banks

with durable informational advantages that should correspond with meaningful changes in bank beliefs. As such, banks may be more likely to incorporate this information into their internal risk models.

There is an alternative channel by which proximity to natural disasters can affect belief formation: salience bias. Banks may overreact to the event by assigning a greater likelihood to similar events of occurring again. In turn, proximity to the event may increase the salience of these effects. We distinguish salience from the informational channel above in the following ways. The effects should be transient such that beliefs should revert to their pre-event levels as the salience of the event wanes over time. In contrast, the relevance of soft information learned from the event should be relatively durable and not diminish over time. Moreover, the effect of the bias should be localized. The occurrence of a natural disaster in one region need not be linked to the likelihood of similar events occurring in other regions. As such, these effects are expected to impact most the areas directly affected as opposed to unaffected areas.

4.2 Bank Learning Tests

Identifying bank beliefs and learning is oftentimes inferred indirectly by examining changes in ex post behavior. However, there are cases where outcomes may not necessarily comport with belief changes, which may suggest alternative mechanisms at work, or where there is an absence of behavioral changes in spite of belief changes. Because our analysis focuses on assessing channels that modulate adaptation behavior, it is critical to observe both.

To overcome these challenges, our empirical strategy centers on extracting bank beliefs from their internal risk models that are directly used in lending decisions. This allows us to directly identify adaptation behavior: namely outcomes that can be directly tied to changes in beliefs.

We take a reduced form approach in recovering weights related to emergent risk factors from bank internal risk models. We regress the loan-level default probabilities on property-

level flood risk measures; the resulting risk loading from the regression should correspond with the weight related to flood risks in the model.

One potential drawback with this approach is that the internal risk models may also include covariates that are unobservable to the econometrician. These factors may include both time-varying credit demand as well as other supply factors. In order to address this issue, we employ high dimensional fixed effects estimators similar to Khwaja and Mian (2008). First, loan fixed effects allow us to restrict the model to time variation within the same household, which mitigates the influence of household and loan characteristics associated with assortative matching factors that may lead to selection bias. Second, interactive fixed effects on the county and date levels allow us to purge any effects related to local demand or regional heterogeneity. Finally, the interactive fixed effects on bank and date levels enable the analysis to focus on intra-bank variation at each point in time to account for potential shocks associated with the hurricane on bank balance sheets.

For loan i, bank b, county g, and date t, the regression model takes the following form:

$$Y_{i,g,b,t} = \beta_0 \cdot FloodRisk_i + \beta_1 \cdot Post_t \times FloodRisk_i + \alpha_i + \delta_{g,t} + \gamma_{b,t} + \varepsilon_{i,g,b,t} \tag{1}$$

where the dependent variable $Y_{i,g,b,t}$ is the natural log of one plus the probability of default for property i (located in county g) for bank b at date t. $FloodRisk_i$ is the natural log of one plus the flood factor for property i. $Post_t$ is a dummy that takes value one if date t occurs after August 2017 and zero otherwise. α_i, $\delta_{g,t}$, and $\gamma_{b,t}$ denote the fixed effects associated with the household, county $\times$ date and bank $\times$ date levels. Robust standard errors are clustered on the household, county-date and bank-date levels.

To examine heterogeneity in bank learning, we use the proportion of bank b's home equity loan portfolio that intersects with the hurricane path as of July 2017, or $BankExposure_b$. We augment the model with interaction terms associated with $BankExposure_b$:

$$Y_{i,g,b,t} = \ldots + \beta_2 \cdot Post_t \times FloodRisk_i \times BankExposure_b + \ldots \tag{2}$$

Our focus will be on the triple interaction term (β_2). A positive coefficient would indicate that banks with higher exposures to the event place greater weight on flood factors following the event relative to banks with lower exposures.

4.3 Adaptation Tests

We next describe the main tests that focus on adaptation behaviors. Those tests examine the changes in lending activities of banks with and without exposures in the aftermath of Hurricane Harvey outside of the affected areas. The data is aggregated to the levels of a bank, a county, and a flood risk group. For each county, we create two risk groups: properties with a flood factor above and below 4. We examine multiple dimensions of bank adaptation by using the following dependent variables:

(a) **Stock changes**: the change in the log total credit limit from July 2017 to September 2019,

(b) **New originations**: the log total credit limit of newly originated loans as of September 2019,

(c) **ARM margins**: the change in the weighted average ARM margin between new originations and existing loans,

(d) **Line actions**: the fraction of loan-months with credit line freezes or reductions in the post-Harvey period.

The estimator takes the following form:

$$\Delta Y_{j,c,k} = \beta \cdot \mathit{HighRisk}_k \times \mathit{BankExposure}_j + \mathbf{X}'_{j,c,k}\boldsymbol{\gamma} + \alpha_j + \delta_{c,k} + \varepsilon_{j,c,k} \tag{3}$$

where j indexes the bank, c the county, and k the risk group. $\mathit{HighRisk}_k$ is an indicator for properties in the high-risk group. Bank fixed effects (α_j) absorb average bank-level factors.

County × risk group fixed effects ($\delta_{c,k}$) absorb local demand conditions specific to each risk level. Standard errors are clustered at the county level. The key coefficient of interest is β, which captures the differential change in outcomes across high- versus low-risk areas for banks with higher versus lower exposures to the natural disaster.

5 Bank Learning Results

This section presents the results from the tests that examine bank learning about flood risks following Hurricane Harvey. Overall, we find evidence that bank exposures to the hurricane is strongly associated with bank learning about future flood risks.

5.1 Flood Risk and Internal Bank Models

We start with simple tests based on univariate regression models where the dependent variable is the loan-level default probability and the explanatory variable is the property-level flood risk measure. Figure 5 displays the marginal effects based on the regression coefficients for each quarterly regression from eight quarters prior through eight quarters following the event. We interpret the marginal effects to correspond with the average weight related to flood risks in the internal risk models. Overlaid on top of the figure are 95% confidence bands. The standard errors are clustered on the ZIP code level. We estimate the marginal effects by estimating the change in the dependent variable given a one-standard deviation increase in the flood risk measure relative to its sample mean. For example, for the +8-quarter subsample, the marginal effect is 1.3%. In comparison, the sample mean of the PD measure is 4.5%.

Several interesting patterns emerge. The effects are not significant until after one quarter following Hurricane Harvey. They suggest that the patterns in the sensitivities for the flood risk measure are associated with responses to the event. It may take some time to develop and incorporate natural disaster risk into internal bank risk models, potentially explaining the

delay in the effects. Importantly, they suggest that the parallel trends assumption required for the validity of the approach does in fact hold.

Additionally, the effect does not reverse over time but rather increases during the first year before levelling off. This suggests that the effect is not transient. That is, banks do not appear to overreact to the event by overweighting the flood factor initially, as found in other studies.

The univariate test results suggest that banks, on average, respond to the natural disaster by updating their internal risk models to account for flood risks. We next turn our attention to the multivariate regression model specifications, which address various endogeneity issues described in the previous section. The results are displayed in Table 3.

We present the results iteratively including additional factors in each specification to help assess their importance. Column (1) only includes loan fixed effects. The estimate on the interaction term is statistically significant at the 1% level and the magnitudes are comparable to the univariate tests. When including time-varying bank and county fixed effects in Column (2), the coefficient remains significant, though attenuates considerably. This suggests that the influence of other factors is quite large, providing some validation to the empirical design. Column (3) decomposes the post-event period by quarter.[2]

5.2 Heterogeneous Bank Learning

We next examine heterogeneity in how internal bank risk models are updated in response to bank exposures to Hurricane Harvey. Table 4 displays the results. Column (1) only includes loan fixed effects. The triple interaction term coefficient is positive and statistically significant at the 1% level. This suggests that the increased weight on flood risks in bank internal risk models are more pronounced for banks with a larger proportion of its loan portfolios in the affected regions. Column (2) also includes time-varying bank and county

[2]The within R^2 for the loan-level regressions is very small (below 0.1% across all specifications), reflecting the fact that the high overall R^2 is driven by the fixed effects rather than the interaction terms of interest. This is expected given the granular nature of the loan-level data and the large number of fixed effects.

fixed effects. The triple interaction remains positive and significant, but also increases in magnitude by more than one-third.

To what extent may the results reflect salience bias instead of information advantages? Figure 6 displays the cost and frequency of billion-dollar disasters over time. The average annual frequency from 1980 to 1999 is 3.3 events, compared to 6.6 events during 2000 to 2014. These events are highly visible, and the associated damages were widely reported in trade journals. This casts at least some doubt that banks without any exposures to the event were unaware of the event. Again, we exclude areas directly affected by the disaster, where we would expect a salience bias to be most pronounced, from the testing sample.

Figure 7 plots out the marginal effects from a model that decomposes the results in Table 4 based on the quarters following the event. Rather than reversing, the patterns indicate that the effects remain positive and stable up through eight quarters following the event. The marginal effect of changes to the bank exposure and flood risk measures through quarter eight on the default probability is 3.9%, which is economically significant compared to the sample mean of 4.5%. These results support the interpretation that the effects are related to informational advantages of the affected banks.

5.3 Robustness Checks

We next consider several robustness checks and evaluate alternative explanations. First, the regression specification uses the natural log transformation of one plus the default probability of the loan rather than that of just the default probability. Chen and Roth (2024) argue issues pertinent to difficulty in the interpretability of results following these types of specifications. We note that 98.8% of the default probabilities used in the analysis are non-zero. We perform the analysis excluding these observations. Table A1 displays the results. In Column (1), the triple interaction term coefficient remains positive and statistically significant. We also consider alternative specifications that address this issue over the full sample. Columns (2) and (3) display the results when replacing the dependent variable for a dummy associated

with the top 25th and 10th sample percentiles for the default probability measure. The triple interaction terms are positive and statistically significant across the two specifications.

We also assess the influence of extreme values of the flood factor measure by considering an alternative specification that maps it to a dummy variable. We find the results remain significant.

Table A3 displays the results when excluding specific states from the analysis (California, Florida, Louisiana, and New York). Across the subsamples, the estimates are all statistically significant at the 1% level.

We directly assess concerns related to enhanced regulatory scrutiny in coastal areas by performing the analysis on sample splits based upon coastal and non-coastal counties. Table A4 displays the results. The coefficients are nearly identical across subsamples.

Finally, we perform the analysis using expected losses as the dependent variable. Table A5 displays the results. Consistent with the main specifications, the triple interaction term coefficient is positive and statistically significant at the 1% level.

5.4 Loan-level Outcomes

We next analyze bank adaptation by checking if the effect of learning corresponds with changes in drawdown behavior and credit line provisioning on existing loans. Table 5 displays the results. Across all specifications, the triple interaction term is statistically significant at the 1% level. Column (1) shows that households in areas with higher flood risks were less able to drawdown on their existing lines of credit following Hurricane Harvey. Column (2) provides a possible explanation, that households in riskier areas received reductions to the credit limit available to them, and that the effect is pronounced for banks with higher exposures.

6 Bank Adaptation Behavior

This section presents the results from tests on bank adaptation outcomes. In the previous section, we provided evidence demonstrating that exposures are associated with heterogeneous learning: banks that were more affected by the hurricane were more likely to be informed about their exposure to future flood risks. Now, we investigate whether informed banks respond by adjusting their HELOC lending in markets with higher flood risk exposure. This may be particularly true for junior-lien property loans like HELOCs, as insurance payouts are less likely to cover losses due to lower seniority in the absolute priority rule. We continue to use loan portfolio exposures to Hurricane Harvey to help us differentiate informed versus uninformed banks with respect to natural disaster risks. As described in Section 4.3, the data is aggregated to the bank × county × risk group level, and the specifications include bank fixed effects and county × risk group fixed effects with clustering at the county level.[3]

6.1 Stock Changes

Table 6 displays the results for changes in the existing HELOC portfolio. The dependent variable is the change in the log total credit limit from July 2017 to September 2019 for bank j in county c and risk group k. The key variable of interest is the interaction between the high-risk indicator and bank exposure. We present three specifications with progressive controls: (1) controlling for the baseline credit limit, (2) adding the baseline utilization rate, and (3) further adding the change in drawn balances.

Across all specifications, the interaction term is negative. Banks with higher exposures to Hurricane Harvey differentially reduce their credit limits in high-risk areas relative to

[3]The adaptation specifications use bank fixed effects and county × risk group fixed effects, which are less demanding than the loan-level learning specifications (loan FE, bank × date FE, county × date FE). The difference reflects the data structure: the adaptation outcomes are measured as long first-differences (July 2017 to September 2019), eliminating the time dimension and leaving a single cross-section at the bank × county × risk level. Bank-by-date fixed effects are not feasible in a single cross-section. The county × risk FE absorbs all local demand variation at each risk level, and the bank FE absorbs bank-level heterogeneity in lending strategy, leaving identification from within-bank, within-county variation across risk groups—the same Khwaja and Mian (2008) logic applied to the cross-section.

low-risk areas, compared to banks with lower exposures. The effect is robust to the inclusion of controls for the utilization rate and changes in drawn balances. These results indicate that the effects of the event on internal bank risk models have tangible consequences for credit provision: informed banks actively contract their HELOC exposure in risky markets.

6.2 New Originations

The results in the previous subsection focused on changes in the existing stock of loans. We next examine whether banks also adjust their new lending activity. Table 7 displays the results. The dependent variable is the log of total credit limits for newly originated loans as of September 2019 that were not in the bank's portfolio as of July 2017.

The interaction term is negative and statistically significant across all specifications. Banks with higher exposures originate fewer new HELOCs in high-risk areas relative to low-risk areas, compared to banks with lower exposures. Importantly, the magnitudes are larger than those for stock changes, suggesting that the adaptation response is more pronounced on the extensive margin of new lending. These results suggest that informed banks primarily adapt by curtailing new credit to risky markets rather than by running off existing loans.

6.3 ARM Margins

Beyond quantities, banks may also adapt through pricing. Table 8 examines this possibility by studying changes in the adjustable-rate mortgage (ARM) margin. The dependent variable is the change in the highlimit-weighted average ARM margin between new originations and the existing loan stock. A positive coefficient would indicate that informed banks charge higher margins in risky areas—consistent with risk-based pricing.

The results indicate that the interaction term is positive and statistically significant. Banks with higher exposures to Hurricane Harvey increase their ARM margins differentially more in high-risk areas. This provides direct evidence that adaptation occurs not only through quantity adjustments but also through pricing: informed banks demand higher

compensation for the flood risk embedded in the collateral.

6.4 Line Actions

Finally, we examine whether banks take active line management actions on existing loans. Table 9 displays the results. The dependent variables are the fraction of loan-months in the post-Harvey period where the credit line was frozen and the fraction where the credit limit was decreased.

The interaction terms are positive and statistically significant for both outcomes. Banks with higher exposures are more likely to freeze credit lines and reduce credit limits in high-risk areas relative to low-risk areas. These results demonstrate that informed banks not only reduce new lending and raise prices, but also actively manage their existing portfolio exposures through administrative actions.

6.5 Pre-Trends and Robustness

A key identification assumption is that informed and uninformed banks had similar trends in adaptation outcomes prior to Hurricane Harvey. While the learning results in Section 5 demonstrate parallel trends in internal risk models (Figure 5), we must verify that pre-existing trends in lending outcomes do not confound the adaptation tests.

We conduct a placebo test by re-estimating the origination specification (Table 7) using pre-Harvey windows as placebos. Bank exposure is held fixed at its July 2017 value, and we construct the same long first-difference over increasingly long pre-Harvey windows. Table A8 reports the results. The six-month window immediately preceding Harvey (January 2017 to July 2017) yields a coefficient of $+0.036$ ($p = 0.86$), confirming that informed and uninformed banks had parallel origination trends in the period directly prior to the event. This contrasts sharply with the actual treatment window (July 2017 to September 2019), where the coefficient is -1.060 ($p < 0.001$). Longer pre-Harvey windows (extending back to 2015–2016) yield significant coefficients, but these periods are confounded by prior flood events

(Louisiana floods of August 2016, Hurricane Matthew in October 2016, Houston flooding in May 2015), oil price fluctuations, and interest rate changes that differentially affected banks with larger Texas and Gulf Coast portfolios.[4]

Additionally, the results are robust to alternative definitions of the high-risk cutoff. Table A9 re-estimates the key specifications using flood factor thresholds of 3, 5, and 6 (compared to the baseline of 4). The origination and peer exposure coefficients maintain the same sign across all cutoffs and remain statistically significant.

Overall, the results in this section provide direct evidence of bank adaptation to natural disaster risks along four distinct dimensions: existing portfolio contraction, reduced new originations, higher pricing, and active line management. A placebo analysis confirms that the effects are not driven by pre-existing trends. These findings suggest that the learning documented in Section 5 translates into concrete changes in bank behavior. The results also suggest that relatively uninformed banks respond by increasing their lending activities in riskier market segments. This raises the possibility that flood risks may become more concentrated on the balance sheets of less informed banks, which may have implications for broader competitive dynamics explored in the next section.

7 Competition and Strategic Complementarities

The previous section provides direct evidence of bank adaptation to future flood risks along multiple dimensions. This section explores the economic mechanisms driving these adaptations, particularly the role of competitive dynamics. We investigate these mechanisms by analyzing variations in bank responses across different market conditions. First, we examine how general competitive conditions influence adaptation behavior on new originations. Second, we assess whether the adaptation behavior of competitors shapes a bank's own lending

[4]The six-month window is chosen as the relevant parallel-trends test because no major flood events occurred during January–July 2017. Oil prices had stabilized at approximately $50/barrel, and the March and June 2017 rate hikes affected all banks symmetrically.

decisions. Finally, we estimate a combined specification that nests both channels.

7.1 Market Concentration

We start by examining the effects of general competitive conditions across local markets on new originations. We augment the baseline specification with the local market concentration ratio (CR3) based on the share of HELOC dollar balances for the top three banks in each county, as well as the Herfindahl-Hirschman Index (HHI). The tests focus on the triple interaction between the risk indicator, bank exposure, and the competition measure. All specifications control for the baseline credit limit.

Table 10 displays the results. The CR3 triple interaction is negative and highly significant across specifications (-0.00136, $p < 0.001$), indicating that the adaptation response on new originations is dampened in markets where the top three lenders collectively hold a larger share. In contrast, the HHI triple interaction, while negative (-0.716), does not achieve statistical significance at conventional levels ($p = 0.158$). These results suggest that competitive pressure from the dominant players—captured by CR3—is the relevant dimension of competition for new lending decisions, rather than overall market fragmentation measured by HHI.

One possible explanation is that banks' decisions about new lending are primarily influenced by the behavior of their closest competitors. When the top three firms control a larger share, the competitive rivalry among these leaders becomes more intense, and exposed banks may feel compelled to keep originating to preserve market position. In contrast, overall market concentration (HHI) reflects the fragmentation of the entire market, which may be less relevant for the forward-looking origination decision.

7.2 Peer Exposure

Next, we examine the role of competitors and their adaptation behavior. To alleviate the potential simultaneity bias, we proxy for peer adaptation by using the average exposure

values of local competitors, i.e., the average Hurricane Harvey exposure of other banks in the same county and risk group. We augment the regression specification with interaction terms with *PeerExposure*, focusing on the three-way interaction between bank exposure, flood risk, and peer exposure. All specifications control for the baseline credit limit.

Table 11 displays the results. The peer exposure triple interaction is negative and significant (-6.359, $p = 0.004$), with effects that are robust to the inclusion of additional controls.

The results provide evidence of a strategic complementarity in adaptation behavior: a bank is more likely to reduce high-risk originations when its competitors are also exposed and therefore also likely to adapt. When competitors lack exposure and are unlikely to adapt, the bank resists adjusting, even if it has learned about the underlying risks. This pattern is consistent with a coordination dynamic in which banks reduce lending in risky market segments only if they anticipate that competitors will do the same—a pattern driven by concerns over market share preservation (Di Maggio et al., 2019).

7.3 Disentangling Concentration and Complementarities

The preceding analysis examines market concentration and peer exposure effects in separate regressions. A natural question is whether these channels are distinct or whether one subsumes the other. Concentration and peer exposure could be correlated—concentrated markets may feature more homogeneous banks with similar exposures—raising the concern that the separate estimates capture overlapping variation. To address this, we estimate a combined specification that simultaneously includes both the concentration triple interaction and the peer exposure triple interaction:

$$\begin{aligned}\ln(1+\mathrm{HL}_{\mathrm{new}})_{jmr} = \alpha_j + \alpha_{mr} + \beta_1 \cdot \mathrm{HR}_r \times \mathrm{BkExp}_j + \beta_2 \cdot \mathrm{HR}_r \times \mathrm{BkExp}_j \times \mathrm{Conc}_m \\ + \beta_3 \cdot \mathrm{HR}_r \times \mathrm{BkExp}_j \times \mathrm{PeerExp}_{jm} + \mathrm{controls} + \varepsilon_{jmr} \qquad (4)\end{aligned}$$

where Conc_m is the concentration measure (CR3 or HHI) and PeerExp_{jm} is the average exposure of competing banks. All lower-order interaction terms are included.

Table 12 displays the results. Two columns present the combined model for originations with CR3 (Column 1) and HHI (Column 2). All specifications include the full set of controls.

Several findings emerge. First, the peer exposure triple interaction remains statistically significant in both columns, at -5.785 ($p = 0.006$) with CR3 and -6.350 ($p = 0.003$) with HHI. These estimates are close to those from the separate peer exposure regressions in Table 11, indicating that the strategic complementarity channel is robust to controlling for market concentration.

Second, the CR3 triple interaction remains significant for originations (-0.0011, $p < 0.001$) in the combined model. The concentration effect identified in Table 10 is therefore not an artifact of omitting peer exposure.

Third, and most notably, the base exposure coefficient (HR $\times$ BkExp) attenuates substantially in the combined model, from -0.918 ($p < 0.001$) in the separate regressions of Table 7 to $+0.407$ ($p = 0.35$). The 95% confidence interval for this attenuated coefficient spans approximately $[-0.4, 1.3]$, reflecting substantial imprecision that prevents us from precisely determining the partial effect of own exposure holding competitive dynamics constant. This imprecision is partly attributable to collinearity between the base interaction and the competition and peer triple terms in the Y-14M sample (variance inflation factor of 36.6 for the base coefficient; see footnote 5).[5] The magnitude of the attenuation—a shift of approximately 1.3 log points from the unconditional estimate—is robust to this collinearity and indicates that the adaptation effect on new originations documented in Section 6 was substantially driven by the competitive dynamics captured by concentration and peer exposure.

[5]The variance inflation factor for HighRisk $\times$ BankExposure in the Y-14M combined specification is 36.6, reflecting correlations of 0.92 with the CR3 triple and 0.90 with the peer triple. This collinearity arises because bank exposure, market concentration, and peer exposure are jointly shaped by geographic market structure among the eight Y-14M banks. In the Equifax sample of 1,118 banks, the VIF for the same coefficient falls to 1.3, confirming that the collinearity is a small-sample phenomenon rather than a structural feature of the model. The Gelbach (2016) decomposition in Section **??** provides a path-independent quantification that does not depend on the precision of individual coefficients in the combined model.

A pooled interaction test in the Equifax sample confirms that the peer and concentration channels are independent: a four-way interaction between the peer triple and market concentration is statistically insignificant ($p = 0.42$), indicating that the two channels capture distinct variation. We develop the implications of these findings in Section 8.

7.4 Robustness

Functional form and the extensive margin. The dependent variable for originations uses the $\ln(1+y)$ transformation, which may be problematic given the 53.5% zero share in the data (Chen and Roth, 2024). Table A11 compares three estimators—OLS on $\ln(1+y)$, Poisson pseudo-maximum likelihood (PPML), and OLS on the inverse hyperbolic sine (IHS)—across the origination, competition, and peer exposure specifications. All three estimators produce coefficients with the same sign and statistical significance. PPML estimates are substantially smaller in magnitude.

To understand this discrepancy, we decompose the treatment effect into its extensive and intensive margins in the Equifax sample (Table A18). The extensive margin—the probability of originating any HELOC in a county-risk cell—accounts for the bulk of the adaptation response. The baseline coefficient on the extensive margin is -0.149 ($p < 0.001$), while the intensive margin conditional on entry is small and statistically insignificant (-0.175, $p = 0.16$). The peer triple follows the same pattern: the extensive margin effect is -0.330 ($p = 0.048$), while the intensive margin is insignificant (-0.846, $p = 0.22$). Banks adapt primarily by ceasing to originate in high-risk areas entirely, not by scaling back within markets where they remain active.

This extensive-margin dominance explains the PPML magnitude discrepancy. PPML weights observations by the conditional mean and therefore downweights the extensive margin relative to OLS on $\ln(1 + y)$. In a setting where the economically relevant adaptation channel is market exit, the OLS estimates capture the first-order effect. The IHS transformation, which accommodates zeros without downweighting the extensive margin, produces

estimates nearly identical to the baseline OLS specification, confirming that the results are not an artifact of the $\ln(1+y)$ functional form. This extensive-margin finding also complements the loan-level evidence in Section 5, which documents intensive-margin adjustments—reduced credit limits, lower drawdowns, and higher pricing—on *existing* loans.

Bank coverage. The peer exposure measure is constructed from the average exposure of competing banks in each county × risk cell. Table A10 reports the distribution of banks per cell: the median cell contains five banks, with 80% of observations coming from cells with four or more banks. We verify that the peer exposure results are robust to restricting the sample to cells with at least three or at least four banks; the coefficients strengthen in the restricted samples (Table A10).[6]

8 A Coordination Framework

Dasol — this is a full revision of Section 8. Key changes at a glance:

1. *The framework (Section 8.1) is replaced. The previous payoff was linear in a, giving corner solutions, two Nash equilibria, and no prediction about the triple interaction. The new payoff is quadratic in a with two complementarity terms: a peer-norm deviation cost (λ_0) and a multiplicative signal–peer term (λ_1) microfounded by adverse-selection spillovers in borrower reapplication. The equilibrium is pinned down by a standard contraction condition (Proposition 1), so Hurricane Harvey shifts a unique equilibrium rather than selecting between several. The global-games apparatus is no longer load-bearing.*

2. *A new Section 8.2 maps the framework to the empirical regressions. The three estimated coefficients identify three structural ratios. Our flagship triple-interaction coefficient β_3*

[6]The Y-14M data cover eight bank holding companies, represented by nine distinct RSSD identifiers due to subsidiary structure. One subsidiary appears only in the loan-level data and not in the county-level aggregation, resulting in nine rather than eight banks in the leave-one-out analysis of Table A6.

is now a direct model prediction, not a loose consistency with the story.

3. *Section 8.3 is rewritten around four qualitative predictions the framework delivers. For this conference version we skip the Gelbach decomposition and the Glaeser–Scheinkman multiplier $\mathcal{M}$. The multiplier formula $\mathcal{M} = 1 + \beta_2/\beta_1$ is clean, but (a) Equifax does not currently report the double interaction HighRisk $\times$ PeerExposure, and (b) reading the ratio off Y-14M goes through a collinearity-driven sign flip of β_1 that invites referee questions. Both quantifications are deferred to the revision draft.*

4. *Sections 8.4 through 8.7 are retained as light edits: the peer-identification tests, the Equifax external-validation and two-way-clustering robustness, the net-credit-supply evidence, and the policy discussion survive with small wording changes to align them with the new framework.*

One item to chase for the next iteration: we need the coefficient on HighRisk $\times$ PeerExposure (without the own-exposure interaction) in the Equifax full spec. That single number closes the multiplier identification in the clean sample. Not for this round.

The preceding sections establish three empirical facts. First, banks exposed to Hurricane Harvey learn about flood risk and adjust their internal risk models (Section 5). Second, informed banks reduce new originations in high-risk areas relative to uninformed banks (Section 6). Third, this adaptation is moderated by local market structure and, crucially, by whether competitors are also likely to adapt (Section 7). We now organize these findings within a deliberately parsimonious framework of strategic coordination among adapting lenders. The framework derives four testable predictions about the cross-sectional patterns of adaptation; the remainder of the section confirms each.

8.1 A framework with coordination

Consider a local lending market of N banks indexed by j. Each bank chooses an adaptation intensity $a_j \in \mathbb{R}_+$, interpreted as the reduction in new high-risk HELOC originations relative

to a pre-Harvey benchmark. Bank j's private signal about flood risk is

$$\theta_j \;\coloneqq\; f(\text{BankExposure}_j) \cdot \text{HighRisk}_c, \tag{5}$$

where $\text{HighRisk}_c \in \{0, 1\}$ indicates whether the cell is high flood risk and $f\colon \mathbb{R}_+ \to \mathbb{R}_+$ satisfies $f' > 0$ and $f'' < 0$. More Harvey exposure strengthens the risk signal; learning exhibits diminishing returns because the first loans in the hurricane zone generate the largest informational gains while subsequent exposure refines the estimate at a decreasing rate. The concavity of f is itself a testable implication, which we confirm below. Peer adaptation in the cell is $\bar{a}_{-j} \coloneqq (N-1)^{-1} \sum_{k \neq j} a_k$.

Bank j's payoff is

$$\pi_j(a_j; \bar{a}_{-j}; \theta_j) \;=\; (\theta_j - c + \lambda_0\, \bar{a}_{-j} + \lambda_1\, \theta_j\, \bar{a}_{-j})\, a_j \;-\; \tfrac{\gamma}{2}\, a_j^2, \tag{6}$$

with $\gamma > 0$, $\lambda_0 \geq 0$, $\lambda_1 \geq 0$, and $c > 0$. Each parameter has a concrete economic interpretation:

- γ is the convexity of adaptation costs. Screening costs rise with the intensity of adaptation because loan-officer time is constrained, model-update and retraining costs are convex, and sharper risk cuts incur rising customer-relationship costs.
- c is the baseline marginal cost of adaptation, capturing foregone margin on originations the bank chooses not to make.
- λ_0 is a *level* complementarity: peer adaptation lowers the level cost of adapting. The microfoundation (Appendix A.2) is a quadratic peer-norm-deviation penalty, justified by reputational and supervisory pressure on banks whose behavior drifts from peer norms.
- λ_1 is a *signal-dependent* complementarity: peer adaptation amplifies the marginal benefit of acting on my own signal. The microfoundation (Appendix A.1) is an adverse-

selection spillover in borrower reapplication. When peers reject risky applicants, those applicants reapply to my bank, and the mass of such rejects scales with the underlying bad-borrower share — precisely the quantity θ_j estimates. This multiplicative structure is what the triple interaction in the regression captures.

The two complementarity parameters arise from non-overlapping mechanisms, and the empirical identification does not require distinguishing them structurally: λ_0 loads on the double interaction HighRisk×PeerExposure, while λ_1 loads on the triple HighRisk×BankExposure×PeerExposure.

The first-order condition of (6) yields the interior best-response function

$$a_j^*(\bar{a}_{-j};\theta_j) \;=\; \frac{1}{\gamma}\,[\,\theta_j(1+\lambda_1\,\bar{a}_{-j}) - c + \lambda_0\,\bar{a}_{-j}\,]_+\,, \tag{7}$$

where $[x]_+ \coloneqq \max\{x,0\}$ encodes the extensive-margin participation threshold. On the interior, a_j^* is linear in $\bar{a}_{-j}$ with slope $(\lambda_0+\lambda_1\theta_j)/\gamma$: banks with stronger signals respond more strongly to peer behavior. This slope heterogeneity is the structural content of the triple-interaction coefficient.

Proposition 1 (Equilibrium uniqueness)**.** *Let* $\boldsymbol{\theta} \coloneqq (\theta_1,\ldots,\theta_N)$*. If*

$$\lambda_0 + \lambda_1 \max_j \theta_j \;<\; \gamma, \tag{8}$$

then the game defined by payoffs (6) *admits a unique Nash equilibrium on* $\mathbb{R}_+^N$*.*

Condition (8) is a contraction condition: strategic complementarity, evaluated at the highest-signal bank, is dominated by cost convexity. The proof (Appendix B) shows that the best-response function is an affine map whose associated matrix has spectral radius below one under (8), so the fixed point is unique. Proposition 1 replaces the two-equilibrium argument in earlier drafts: Hurricane Harvey *shifts* a unique equilibrium rather than *selecting* between multiple equilibria. The global-games apparatus of Morris and Shin (1998) and Goldstein

and Pauzner (2005) is no longer load-bearing; uniqueness holds even when banks observe noisy private signals of peer exposure, under the same contraction condition.

8.2 Mapping to the reduced-form regression

The empirical specifications of Tables 11, 12, and 13 estimate

$$\begin{aligned} \ln\!\left(1+\text{Limit}_{jc}^{\text{new}}\right) &= \beta_1\,(\text{HighRisk}_c \cdot \text{BankExposure}_j) \\ &+ \beta_2\,(\text{HighRisk}_c \cdot \text{PeerExposure}_{-j,c}) \\ &+ \beta_3\,(\text{HighRisk}_c \cdot \text{BankExposure}_j \cdot \text{PeerExposure}_{-j,c}) \;+\; \text{FE}_{jc} + \varepsilon_{jc}, \end{aligned} \tag{9}$$

where the dependent variable is log new-origination volume in cell c for bank j. Let $\phi \coloneqq f'(\overline{\text{BankExposure}})$ denote the local slope of the signal function at the sample mean. A first-order expansion of the symmetric equilibrium of (6) around $\overline{\text{BankExposure}}$, combined with the equilibrium fixed-point restriction that $\bar{a}_{-j}$ equals peers' equilibrium adaptation (Appendix C), yields the coefficient dictionary

$$\beta_1 \;=\; -\frac{\phi}{\gamma}, \qquad \beta_2 \;=\; -\frac{\phi\,\lambda_0}{\gamma(\gamma-\lambda_0)}, \qquad \beta_3 \;=\; -\frac{\phi^2\,\lambda_1}{\gamma(\gamma-\lambda_0)}. \tag{10}$$

The sign flip relative to (7) reflects that adaptation a_j corresponds to a *reduction* in originations, so the regression coefficients carry the opposite sign.

Three implications of the dictionary are worth highlighting. First, the triple-interaction coefficient β_3 is a direct structural prediction: the cross-partial $\partial^2 a_j^*/(\partial\theta_j\,\partial\bar{a}_{-j}) = \lambda_1/\gamma > 0$ is strictly positive on the interior, which delivers $\beta_3 < 0$ in the originations regression. Second, the full specification that includes the peer triples is required for structural identification of β_1: the short regression omitting the peer terms absorbs part of the equilibrium feedback into the own-exposure coefficient, so the short-spec β_1 does not recover the direct effect $-\phi/\gamma$. Third, the ratio $\beta_3/\beta_2 = \phi\,\lambda_1/\lambda_0$ is a combination of deep parameters that should

be sample-invariant under correct specification; its stability across the Y-14M and Equifax samples is a specification check.

For this version we stop at the qualitative mapping. A structural multiplier quantification — $\mathcal{M} = 1/(1-\lambda_0/\gamma) = 1+\beta_2/\beta_1$ in the Glaeser and Scheinkman (2003) tradition — is a natural further step that we defer. Clean implementation of $\mathcal{M}$ requires the coefficient on HighRisk $\times$ PeerExposure in the Equifax full specification, which is not reported in Table 13; the Y-14M estimate is available but is distorted by collinearity between own and peer exposure (variance inflation factor ≈ 36.6 in the combined specification), which flips the sign of β_1 and makes the ratio imprecise. We return to $\mathcal{M}$ in the revision.

8.3 Empirical patterns consistent with the framework

The framework delivers four qualitative predictions beyond the baseline that exposed banks adapt more. Each is verified in the data; the structural logic of the test is summarized before the empirical number.

(P1) Concavity in own exposure. Because $\theta_j = f(\text{BankExposure}_j)$ with $f'' < 0$, the adaptation response is concave in BankExposure. The concavity test distinguishes the information channel from alternatives such as portfolio rebalancing or regulatory pressure, which predict linear or convex responses. The concavity specification (reported in Section 7) confirms: the quadratic own-exposure term in the Equifax sample is positive and significant ($p = 0.009$), and remains so with peer controls ($p = 0.011$), indicating diminishing marginal adaptation to additional exposure.

(P2) Negative signal–peer interaction. The framework predicts $\beta_3 < 0$ through $\lambda_1/\gamma > 0$. Table 11 reports $\widehat{\beta}_3 = -6.194$ ($p < 0.01$) in the full Y-14M specification (Column 3). Table 13 reports $\widehat{\beta}_3 = -4.252$ ($p = 0.033$) in the peer-only Equifax specification (Column 4) and $\widehat{\beta}_3 = -4.264$ ($p = 0.032$) in the combined Equifax specification with both HHI and

peer channels (Column 6). Sign and significance are stable across specifications and across samples separated by two orders of magnitude in the number of banks (8 vs. 1,118).

(P3) Sign reversal for non-local peers. The adverse-selection spillover that micro-founds $\lambda_1\theta_j\bar{a}_{-j}$ is spatially local: rejected borrowers reapply within the same local market, not across markets. Non-local peer exposure should not enter the payoff at all. A correlated-local-demand alternative, by contrast, would predict the same sign for local and non-local peer interactions, since both pick up the same aggregate shock. The observed loading on the non-local peer triple is *opposite* in sign to the local peer triple: $+6.944$ ($p = 0.005$) versus -5.667 ($p = 0.011$) in the non-local placebo specification (Section 8.4). The sign reversal is inconsistent with correlated-shock explanations and directly consistent with the local-spillover structure of the framework.

(P4) Attenuation of the own-exposure coefficient in the full specification. The framework predicts that β_1 recovers the direct effect $-\phi/\gamma$ only in the full specification; the short specification absorbs equilibrium feedback into the own-exposure coefficient, so the two estimates will differ. In Y-14M the attenuation is substantial: $\widehat{\beta}_1$ moves from -0.918 ($p < 0.001$) in Table 7, Column 3 (short specification) to $+0.407$ ($p = 0.35$) in Table 12, Column 1 (full specification with CR3 + peer) and $+0.891$ ($p = 0.07$) in Column 2 (full specification with HHI + peer). The change of roughly one log point is economically large and points to competitive dynamics as a first-order determinant of the observed contraction. The sign flip of $\widehat{\beta}_1$ across specifications reflects the high collinearity between own and peer exposure in the eight-bank Y-14M sample — the ratio of coefficients, rather than any individual coefficient, is the structurally meaningful object. In the Equifax sample with 1,118 banks and substantially lower collinearity, the pattern differs: $\widehat{\beta}_1$ becomes *more* negative when the peer channel is added, moving from -1.819 in Column 1 to -2.580 in the combined Column 6 of Table 13. We interpret the opposite directionality as consistent with differences in the joint distribution of own and peer exposure across the two samples, not as contradictory evidence

on the framework: in both samples the own-exposure coefficient moves substantially when the peer channel enters, and the peer triple is negative and significant.

8.4 Identifying the peer effect

The peer-exposure triple interaction is central to the framework, and its interpretation as strategic interaction rather than correlated local demand requires a defensible identification argument. A natural concern is that peer exposure mechanically shares variation with own exposure because banks lending in the same county have overlapping Harvey footprints. Five tests address this concern.

Permutation test. We randomly reassign bank identities within each county–risk cell while preserving the market structure, recompute peer exposure from the shuffled assignments, and re-estimate the triple interaction. Across 1,000 permutations, the actual peer coefficient falls at the 0.5th percentile of the permutation distribution ($p = 0.005$), rejecting the null that the peer effect reflects market-level sorting. The independent 1,118-bank Equifax sample yields a consistent result ($p = 0.014$).

Non-local peer placebo. If the peer effect reflects local strategic interaction (as the framework's adverse-selection microfoundation requires), peer exposure constructed from the same banks' activities in *other* counties should have no effect — or carry the opposite sign, reflecting competitive substitution rather than complementarity. The local peer triple is -5.667 ($p = 0.011$): when local competitors are exposed, the bank pulls back. The non-local peer triple is $+6.944$ ($p = 0.005$): when the same competitors are exposed through their portfolios in *other* counties, the bank lends *more*. A correlated-shock alternative would predict the same sign in both cases; the observed sign reversal is direct evidence that the peer coefficient reflects local strategic interaction.

Aggregate exposure control. We add county-level aggregate bank exposure as a direct control, so that the peer effect is identified from variation in the *composition* of local competitors rather than the *level* of aggregate local exposure. The peer triple strengthens to -9.734 ($p < 0.001$), confirming that it reflects the identity of specific competitors, not aggregate local conditions.

Leave-one-out. We drop each of the eight Y-14M banks in turn and re-estimate the peer and origination specifications. All leave-one-out estimates retain the same sign (9/9 for both originations and peer exposure), and the majority remain statistically significant (7/9 for each). The peer-exposure coefficient ranges from -1.13 to -11.47 across subsamples, reflecting the sensitivity of the estimate to individual banks' contributions in the eight-bank sample. The Equifax replication (-4.252, $p = 0.033$) confirms that the result is not driven by any single Y-14M bank.

Capacity reallocation. A further concern is that the peer-exposure effect reflects capacity reallocation rather than strategic interaction: when competitors are exposed to Hurricane Harvey in *other* markets, they may shift lending capacity across counties, changing competitive conditions independently of any coordination motive. We test this directly in the Equifax sample, where the large cross-section ensures that peer exposure and competitor capacity are essentially uncorrelated ($r = -0.01$). We construct competitor capacity as the average change in competitors' lending volume in all other counties and add it as a triple interaction alongside peer exposure. Controlling for competitor capacity leaves the peer triple unchanged (-4.36, $p = 0.036$) while the capacity measure itself is insignificant ($p = 0.94$). This rules out capacity reallocation as an explanation for the peer effect.

8.5 External validation and inference

The FR Y-14M data deliver precise identification but cover only the eight largest bank holding companies, which limits the asymptotic validity of clustered inference. To address

both external validity and inference, we replicate the analysis using credit-bureau data from Equifax covering 1,118 bank holding companies matched to FR Y-9C regulatory filings.

Table 13 reports the results. The baseline adaptation result replicates: $\widehat{\beta}_1 = -1.819$ ($p < 0.001$; Column 1). Market concentration as measured by HHI dampens the adaptation response (1.307, $p = 0.018$; Column 3), consistent with Table 10. The peer-exposure triple is -4.252 ($p = 0.033$; Column 4), the same sign and comparable magnitude as the Y-14M estimate. In the combined specification (Column 6), the peer triple retains its magnitude and significance (-4.264, $p = 0.032$) while the HHI triple is absorbed, mirroring the Y-14M pattern in Table 12.

To assess inference robustness, we re-estimate the Equifax specifications with two-way clustering at the county and bank levels. With 1,118 bank clusters the asymptotic approximation for clustered inference is well-founded. The baseline own-exposure coefficient remains highly significant ($p < 0.001$); the peer triple retains significance ($p = 0.039$); the combined-specification peer coefficient is also robust ($p = 0.039$). Standard-error ratios between two-way and county-only clustering range from 0.91 to 1.31, confirming that the additional bank-level cluster dimension does not qualitatively change inference.[7]

8.6 Net credit supply effects

A natural question is whether uninformed banks fill the gap left by informed-bank pullbacks, and whether the peer channel operates independently of exposure status. The framework is silent on this by default — the learning signal θ_j differentiates banks by Harvey exposure — but the coordination channel $\lambda_0 \bar{a}_{-j}$ should operate on all banks in the cell because the peer-norm deviation cost is exposure-invariant.

Results from the Equifax sample, where the large cross-section provides adequate power, confirm this. The coefficient on HighRisk $\times$ BankExposure is -0.964 ($p = 0.008$), indicating

[7]The Y-14M data contain only eight bank clusters, too few for reliable asymptotic inference under two-way clustering. We therefore use the Equifax two-way clustering results as the relevant robustness check for inference.

that all banks — both exposed and unexposed — reduce high-risk originations in proportion to their exposure. The interaction HighRisk×Exposed is -1.004 ($p = 0.035$), confirming that exposed banks pull back significantly more, consistent with the additional learning signal $\theta_j > 0$ in the framework. In the combined specification that includes both own-exposure and peer-exposure interactions with exposure status, the peer effect remains significant (-7.018, $p = 0.016$) and is robust to controlling for whether the bank itself is exposed. The coordination channel operates on top of individual learning, just as the framework predicts. The net effect on total credit supply in high-risk areas is -2.61 ($p < 0.001$). Uninformed banks do not fill the gap, so the coordinated adaptation has real consequences for credit access in high-risk areas, not merely a reallocation across bank types.

8.7 Policy implications

The coordination framework yields a specific policy prediction. If the binding constraint on adaptation is not individual awareness of risk but uncertainty about whether competitors will also adapt, a natural policy response is one that reduces this uncertainty.

Current climate stress-testing exercises, conducted by the Bank of England, the European Central Bank, and the Federal Reserve, evaluate individual bank responses to climate scenarios in isolation. Our results suggest that such exercises may mischaracterize the systemic response because they do not capture the coordination channel: a bank evaluated in isolation exhibits an attenuated adaptation response, while the full equilibrium adaptation emerges only when banks anticipate that competitors facing the same scenario will also adjust.

The framework suggests a complementary tool: climate risk disclosure. Standardized reporting of climate risk exposures — such as those proposed under the Task Force on Climate-related Financial Disclosures and related initiatives — would make risk information common knowledge across market participants. In the framework, disclosure raises the predictability of $\bar{a}_{-j}$ by revealing that competitors face similar risks, thereby enabling the coordination equilibrium without a regulatory mandate on specific risk management prac-

tices. Hurricane Harvey served this function accidentally by simultaneously exposing multiple banks in the same markets to a common shock that made peer adaptation predictable. A disclosure regime would institutionalize the mechanism, allowing coordination to occur in response to the gradual emergence of climate risks rather than requiring a catastrophic event to trigger collective action. This is a distinct policy channel from microprudential interventions (capital charges on individual bank exposures) and conventional macroprudential tools (system-wide stress tests with feedback effects). The coordination mechanism operates through *information*, not incentives or constraints.

Welfare considerations. Coordinated adaptation may improve bank stability by reducing aggregate flood-risk exposure, but may also restrict credit access for borrowers in high-risk areas. If banks are correcting a previous underpricing of flood risk, the pullback is efficient. If, however, banks overreact — for example because the complementarity λ_1 amplifies the response beyond what individual risk assessment would warrant — the result is a credit crunch that depresses property values and disproportionately affects households in vulnerable areas. The welfare implications depend on whether the coordinated response converges to the socially optimal level of risk management or overshoots it. Our framework does not identify the socially optimal level, as we do not observe the true cost of flood risk. The net credit-supply decline documented above — where uninformed banks fail to offset the pullback — suggests that coordinated adaptation has distributional consequences that warrant attention in the design of disclosure policies.

General equilibrium and spatial dynamics. Our analysis focuses on partial-equilibrium effects. In general equilibrium, the credit-supply contraction may feed back into property values, collateral values, and further bank risk through capitalization effects (Kuminoff, Parmeter, and Pope 2013; Hsiang, Oliva, and Walker 2019). If coordinated bank pullback lowers property values in high-risk areas, this reduces collateral values and raises loan-to-value ratios, potentially *increasing* rather than decreasing bank risk exposure. Bakkensen

and Barrage (2022) and Keys and Mulder (2022) document that climate-risk capitalization into property values is incomplete and heterogeneous across markets, suggesting that such feedback effects are plausible. A full general-equilibrium analysis of these spatial dynamics is beyond the scope of this paper, but the partial-equilibrium effects we document represent the first-order channel through which climate risk propagates through the banking system.

Scope and generalizability. Our analysis focuses on HELOCs, which banks retain on their balance sheets because of the absence of a liquid secondary market. The coordination dynamics may differ in markets where originators can securitize risk: in the originate-to-distribute model, banks offload flood-risk exposure, muting both the private signal θ_j and the incentive to adapt. Strategic complementarities may therefore be weaker in securitized markets — not because coordination forces are absent, but because the individual learning channel that initiates the coordination cascade is muted. Our setting isolates the coordination mechanism precisely because the absence of securitization forces banks to internalize risk. Our identification relies on a single event (Hurricane Harvey), raising the question of whether the dynamics generalize. The mechanism — competitive dynamics moderating individual adaptation when information is heterogeneous — should operate whenever a common shock provides asymmetric information to banks with overlapping markets. The key scope conditions are that the shock must be (i) informative about risk, (ii) publicly observable (enabling inference about peer signals), and (iii) geographically heterogeneous in exposure.

9 Conclusion

This paper studies how competitive dynamics shape bank adaptation to emergent risks. Using confidential supervisory data joined to granular flood risk measures, we document that banks exposed to Hurricane Harvey learn about future flood risks and subsequently reduce new lending to high-risk areas. However, this adaptation is not primarily an individual risk management response. In a combined specification that controls for market concentration

and peer exposure, the individual incentive to adapt attenuates substantially—though the precise magnitude of the partial effect remains uncertain due to collinearity in the supervisory sample. A formal Gelbach (2016) decomposition confirms that peer exposure accounts for the bulk of the attenuation (bootstrap 95% CI: [84%, 168%]). Competitive dynamics—market concentration, which raises the cost of unilateral adaptation, and peer exposure, which conditions adaptation on competitors' expected behavior—are first-order determinants of the observed credit contraction. The adaptation response on new originations operates primarily through the extensive margin, with banks ceasing to originate in high-risk areas entirely. On existing loans, the supervisory data reveal intensive-margin adjustments: reduced credit limits, lower drawdowns, and higher pricing. The adaptation response is concave in bank exposure ($p = 0.009$), consistent with a learning mechanism in which the informational value of disaster exposure exhibits diminishing returns.

We subject the peer effect to extensive identification tests. A permutation test rejects the hypothesis that it reflects common local shocks ($p = 0.005$ in supervisory data; $p = 0.014$ in credit bureau data). In an independent sample of over one thousand bank holding companies from credit bureau data, local peer exposure reduces originations while non-local peer exposure *increases* them—a sign reversal that rules out correlated shocks and confirms local strategic interaction. The peer effect strengthens when controlling for county-level aggregate exposure, and the results are robust to two-way clustering, alternative risk cutoffs, Poisson pseudo-maximum likelihood estimation, and leave-one-out bank subsamples. A placebo test confirms parallel pre-trends in adaptation outcomes.

We develop an interpretive coordination framework in which informed banks are near the margin between adapting and not adapting, and a common information shock—Hurricane Harvey—resolves the coordination problem by making peer adaptation more predictable. Exposed banks reduce lending significantly more than unexposed banks, but the peer effect is robust to controlling for exposure status, indicating that the coordination channel operates on top of individual learning. The coordinated pullback has real consequences: uninformed

banks do not fill the gap, and total credit supply in high-risk areas declines significantly.

Our findings have three implications. First, climate stress tests that evaluate banks individually will find an attenuated adaptation response, because the individual incentive to adapt, net of competitive dynamics, attenuates substantially; the full adaptation emerges in equilibrium when banks anticipate that competitors facing the same scenario will also adjust. Second, climate risk disclosure—by making risk information more widely available—can complement regulatory mandates by making peer adaptation predictable, enabling collective adaptation without requiring a catastrophic event to serve as a coordination device. Third, while coordinated adaptation may improve bank stability, it may also restrict credit access for borrowers in vulnerable areas. The design of disclosure policies should account for these distributional consequences.

Our analysis focuses on HELOCs—a non-securitized product where banks retain risk—and exploits a single natural disaster. The coordination mechanism should operate more broadly whenever a common shock provides asymmetric information to competing lenders, but it may be weaker in markets where originators can offload risk through securitization. The development of richer theoretical frameworks that endogenize the coordination threshold, and the identification of settings that enable tests of cross-sectional predictions, remain promising directions for future work.

More broadly, the absence of regulatory guidance on emergent risks does not merely leave banks to manage risks independently; it creates a coordination problem in which the equilibrium level of risk management depends on competitive dynamics rather than individual risk awareness.

References

Alok, S., N. Kumar, and R. Wermers (2020). Do fund managers misestimate climatic disaster risk. *The Review of Financial Studies 33*(3), 1146–1183.

Bakkensen, L., T. Phan, and T.-N. Wong (2025). Leveraging the disagreement on climate change: Theory and evidence. *Journal of Political Economy 133*(10), 3132–3166.

Bakkensen, L. A. and L. Barrage (2022). Going underwater? *The Review of Financial Studies 35*(8), 3666–3709.

Baldauf, M., L. Garlappi, and C. Yannelis (2020). Does climate change affect real estate prices? *The Review of Financial Studies 33*(3), 1256–1295.

Bernstein, A., M. T. Gustafson, and R. Lewis (2019). Disaster on the horizon. *Journal of Financial Economics 134*(2), 253–272.

Blickle, K., E. Perry, and J. a. A. Santos (2024). Do mortgage lenders respond to flood risk? FRB of New York Staff Report No. 1101.

Brunetti, C. et al. (2024). Climate-related financial stability risks. Federal Reserve Bank of New York Economic Policy Review.

Bryant, J. (1980). A model of reserves, bank runs, and deposit insurance. *Journal of Banking and Finance 4*(4), 335–344.

Chen, J. and J. Roth (2024). Logs with zeros? Working Paper.

Correa, R., A. He, C. Herpfer, and U. Lel (2022). The rising tide lifts some interest rates. Working Paper.

Dessaint, O. and A. Matray (2017). Do managers overreact to salient risks? *Journal of Financial Economics 126*(1), 97–121.

Di Maggio, M., A. Kermani, and S. Korgaonkar (2019). Partial deregulation and competition. *Management Science 65*(10), 4676–4711.

Diamond, D. W. and P. H. Dybvig (1983). Bank runs, deposit insurance, and liquidity. *Journal of Political Economy 91*(3), 401–419.

Drechsler, I., A. Savov, and P. Schnabl (2017). The deposits channel of monetary policy. *The Quarterly Journal of Economics 132*(4), 1819–1876.

Egan, M., A. Hortaçsu, and G. Matvos (2017). Deposit competition and financial fragility. *American Economic Review 107*(1), 169–216.

Financial Crisis Inquiry Commission (2011). *The Financial Crisis Inquiry Report, Authorized Edition: Final Report of the National Commission on the Causes of the Financial and Economic Crisis in the Un*. PublicAffairs.

Gallagher, J. and D. Hartley (2017). Household finance after a natural disaster. *American Economic Journal: Economic Policy 9*(3), 199–228.

Gelbach, J. B. (2016). When do covariates matter? and which ones, and how much? *Journal of Labor Economics 34*(2), 509–543.

Goldstein, I. and A. Pauzner (2005). Demand–deposit contracts and the probability of bank runs. *The Journal of Finance 60*(3), 1293–1327.

Keys, B. J., T. Mukherjee, A. Seru, and V. Vig (2009). Financial regulation and securitization. *Journal of Monetary Economics 56*(5), 700–720.

Keys, B. J. and P. Mulder (2022). Neglected no more. Working Paper.

Khwaja, A. I. and A. Mian (2008). Tracing the impact of bank liquidity shocks. *American Economic Review 98*(4), 1413–1442.

Kousky, C. (2010). Learning from extreme events. *Land Economics 86*(3), 395–422.

Kozlowski, J., L. Veldkamp, and V. Venkateswaran (2020). The tail that wags the economy. *Journal of Political Economy 128*(8), 2839–2879.

Meisenzahl, R. (2023). How climate change shapes bank lending. FRB Chicago Working Paper.

Morris, S. and H. S. Shin (1998). Unique equilibrium in a model of self-fulfilling currency attacks. *American Economic Review*, 587–597.

Murfin, J. and M. Spiegel (2020). Is the risk of sea level rise capitalized in residential real estate? *The Review of Financial Studies 33*(3), 1217–1255.

NOAA (2018). Costliest u.s. tropical cyclones tables updated.

OCC, Federal Reserve, FDIC (2023). Principles for climate-related financial risk management.

Ouazad, A. (2022). Do investors hedge against green swans? Working Paper.

Ouazad, A. and M. E. Kahn (2021). Mortgage finance and climate change. *Review of Financial Studies 35*(8), 3617–3665.

Risser, M. D. and M. F. Wehner (2017). Attributable human-induced changes in the likelihood and magnitude of the observed extreme precipitation during hurricane harvey. *Geophysical Research Letters 44*(24).

Rochet, J.-C. and X. Vives (2004). Coordination failures and the lender of last resort. *Journal of the European Economic Association 2*(6), 1116–1147.

Trenberth, K. E. et al. (2018). Hurricane harvey links to ocean heat content and climate change adaptation. *Earth's Future*.

Vives, X. (2014). Strategic complementarity, fragility, and regulation. *The Review of Financial Studies 27*(12), 3547–3592.

Wang, Y., T. M. Whited, Y. Wu, and K. Xiao (2022). Bank market power and monetary policy transmission. *The Journal of Finance 77*(4), 2093–2141.

Weitzman, M. L. (2009). On modeling and interpreting the economics of catastrophic climate change. *The Review of Economics and Statistics 91*(1), 1–19.

Whited, T. M., Y. Wu, and K. Xiao (2021). Low interest rates and risk incentives for banks with market power. *Journal of Monetary Economics 121*, 155–174.

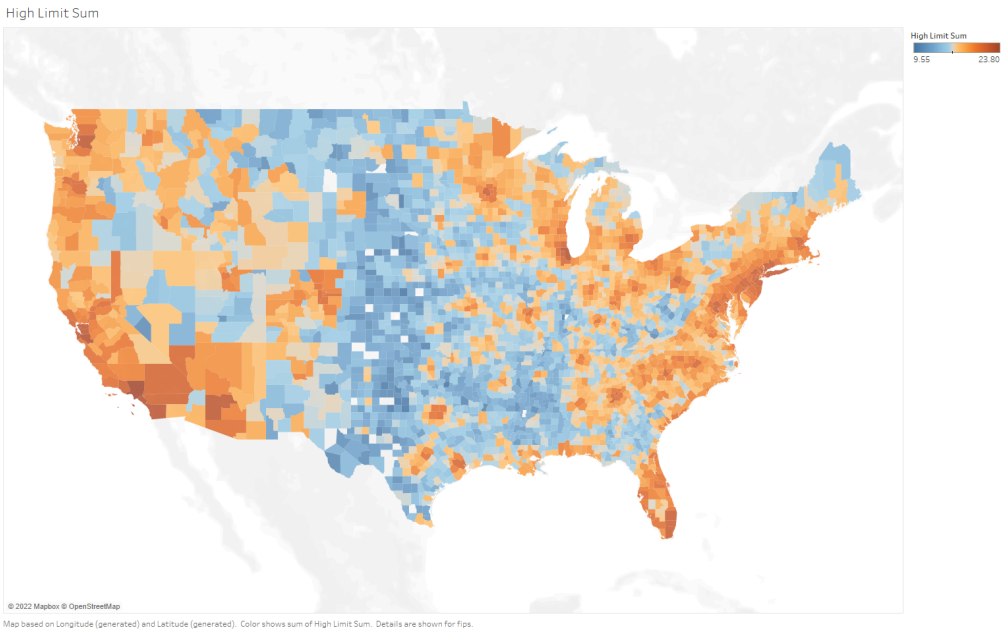
Zhang, W. et al. (2018). Increasing flood risk associated with global warming.

Figure 1: Loan Coverage by County

The figure displays maps of the log of home equity loans by county. Panel (a) is based on the regulatory bank data, and Panel (b) is based on the credit bureau data.

Sources: FR Y-14M (Schedule B.1), Equifax, and Authors' analysis.

(a) Regulatory Bank Data



(b) Credit Bureau Data

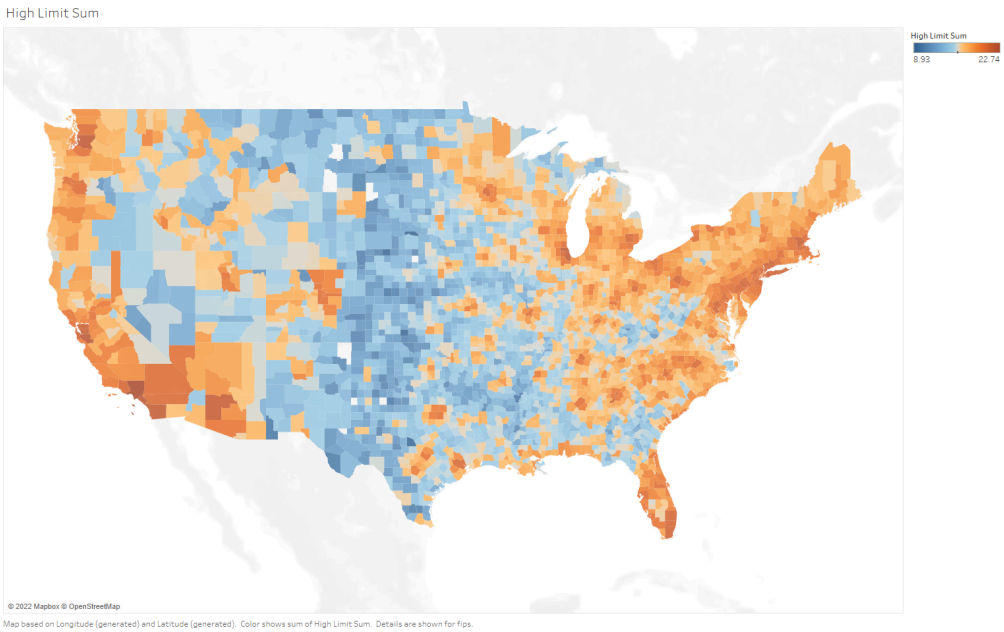


Figure 2: Areas Affected by Hurricane Harvey

The figure displays a map of the trajectory of Hurricane Harvey.

Sources: NOAA and Authors' analysis.

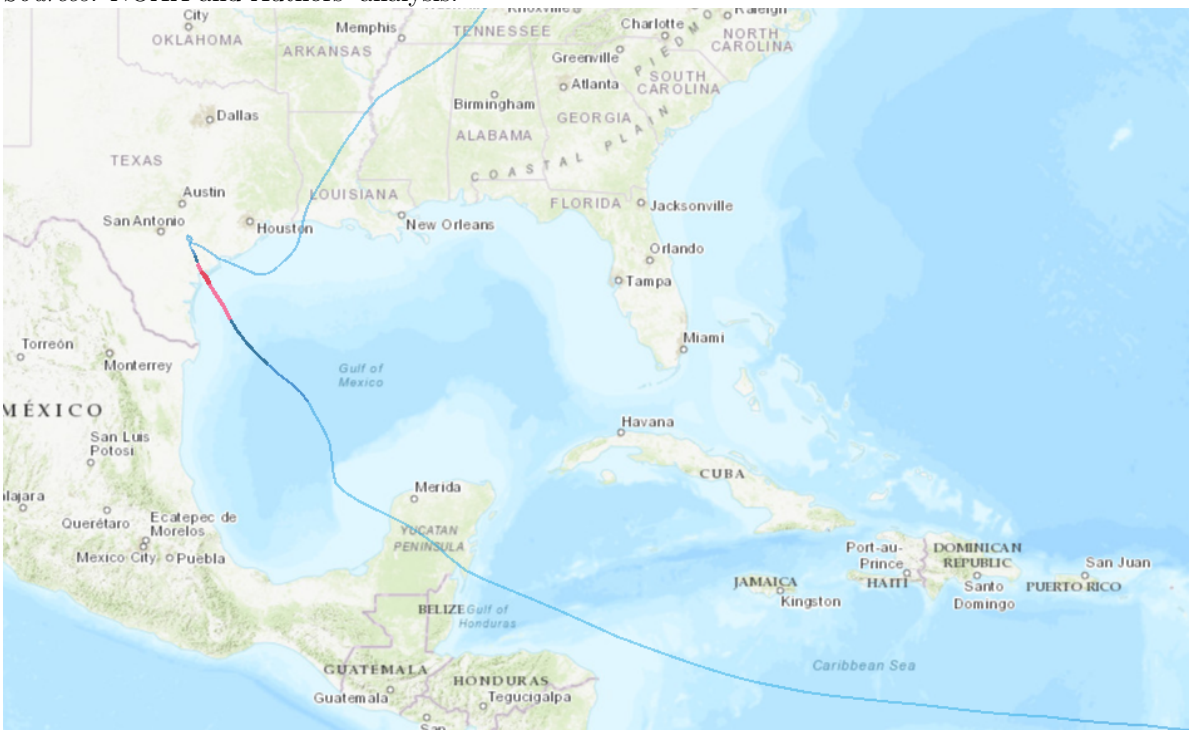


Figure 3: Geographic Distribution of Flood Risks

The figure displays a map of log of the average FloodRisk using the publicly available zip code level data from the First Street Foundation.

Sources: First Street Foundation and Authors' analysis.

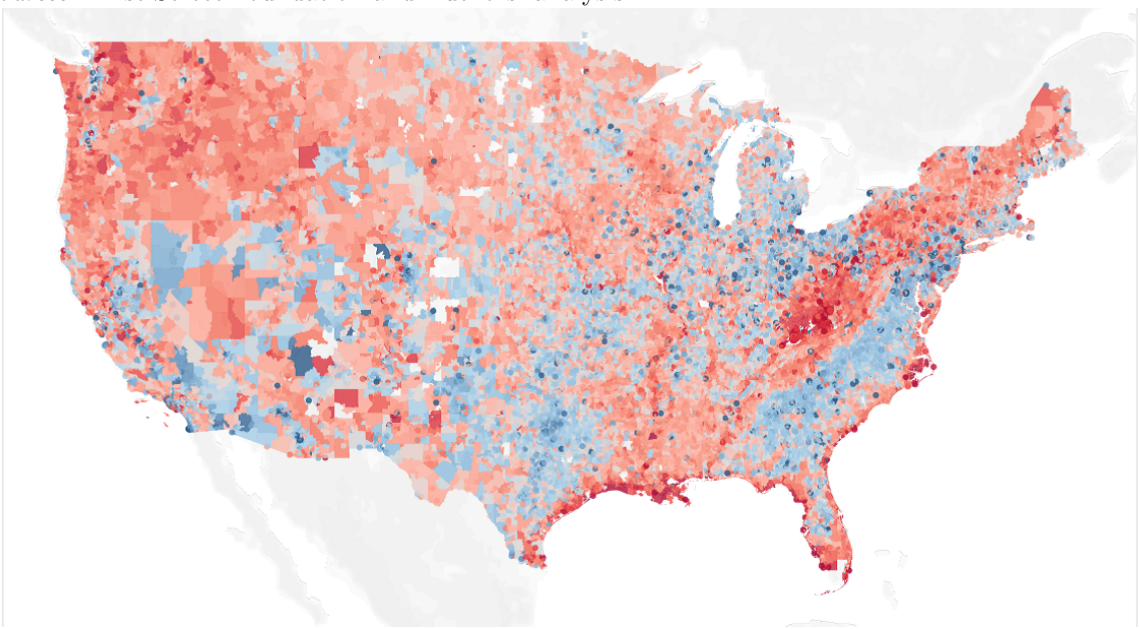


Figure 4: Heterogeneity in Bank Exposures to Hurricane Harvey

The figure displays a map of counties to indicate variation of BankExposure across banks within counties in which banks lend. The choropleth is configured such that counties where all banks have no exposure are colored light grey, where there are both banks with and without exposures are colored dark grey, and where all banks with some exposure are colored orange.

Sources: FR Y-14M (Schedule B.1) and Authors' analysis.

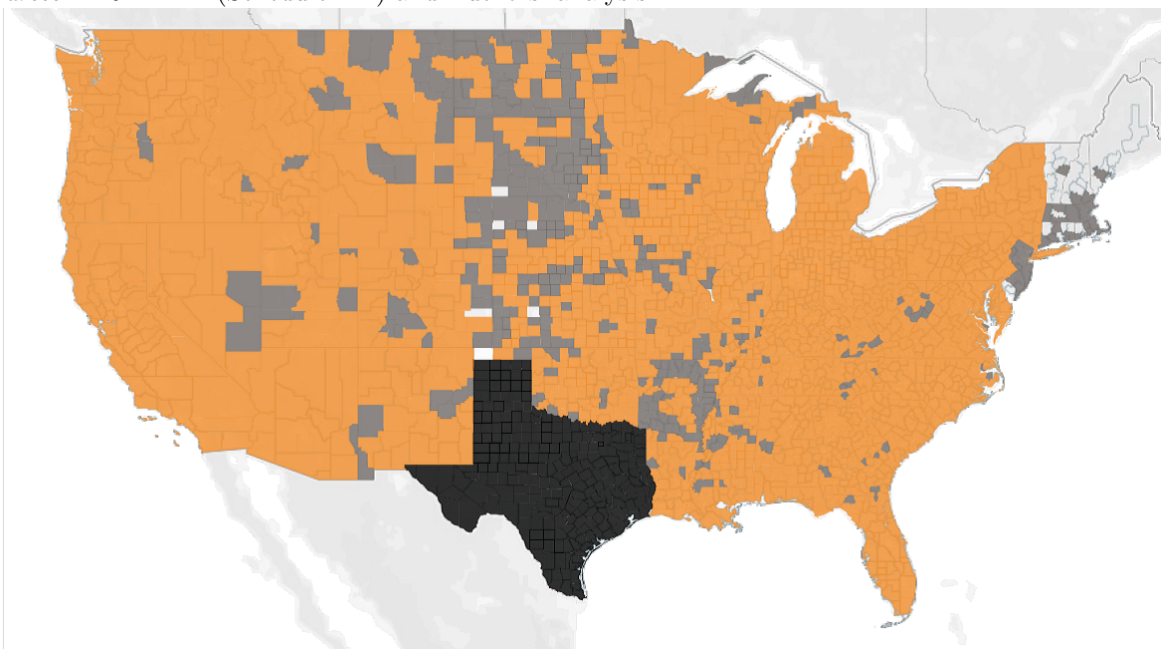


Figure 5: Marginal Effects of Flood Risk on Default Probabilities Around Natural Disaster

The figure displays calculations from univariate regression coefficients where the dependent variable is the banks' estimates of probability of default (PD) for each borrower and the explanatory variable is FloodRisk. The figure shows the estimated average marginal effects with 95% confidence bands by quarter surrounding the event.

Sources: FR Y-14M (Schedule B.1), First Street Foundation, and Authors' analysis.

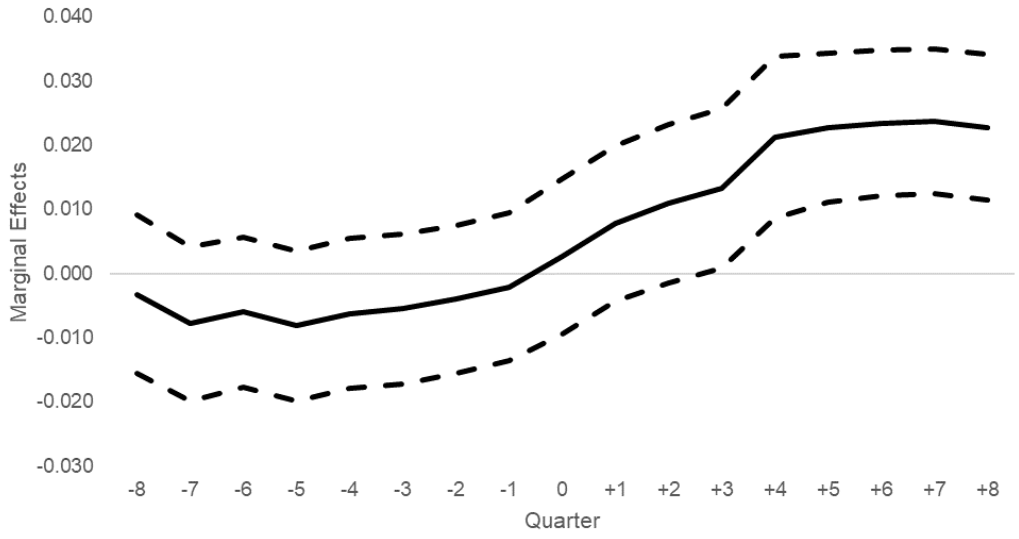


Figure 6: Frequency and Costs of Billion Dollar Disasters over Time

The figure displays the annual frequency and total costs of storms that cause at least $1 billion in damages, with the damage amount adjusted by the CPI.

Sources: NOAA/FEMA and Authors' analysis.

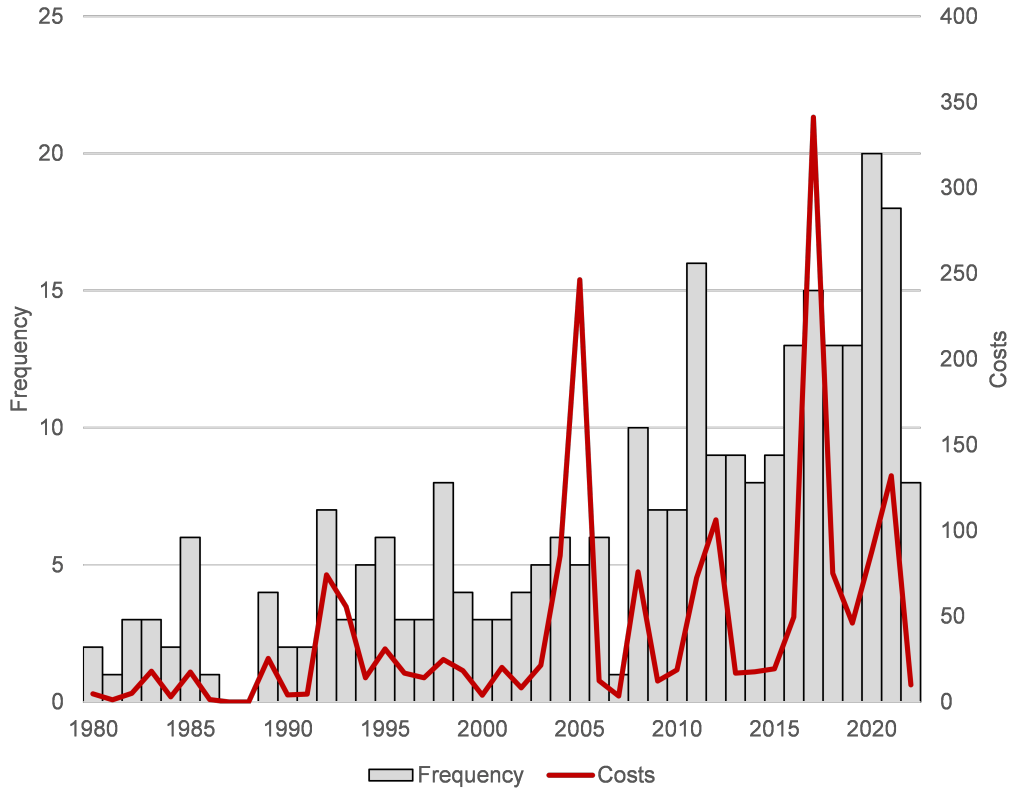


Figure 7: Marginal Effects of Bank Exposures and Flood Risk

The figure displays a figure of regression analysis results where the dependent variable is the banks' estimates of probability of default (PD) for each borrower and the explanatory variables are interaction terms between quarter dummies and FloodRisk. The figure shows the estimated average marginal effects with 95% confidence bands by quarter surrounding the event.

Sources: FR Y-14M (Schedule B.1), First Street Foundation, and Authors' analysis.

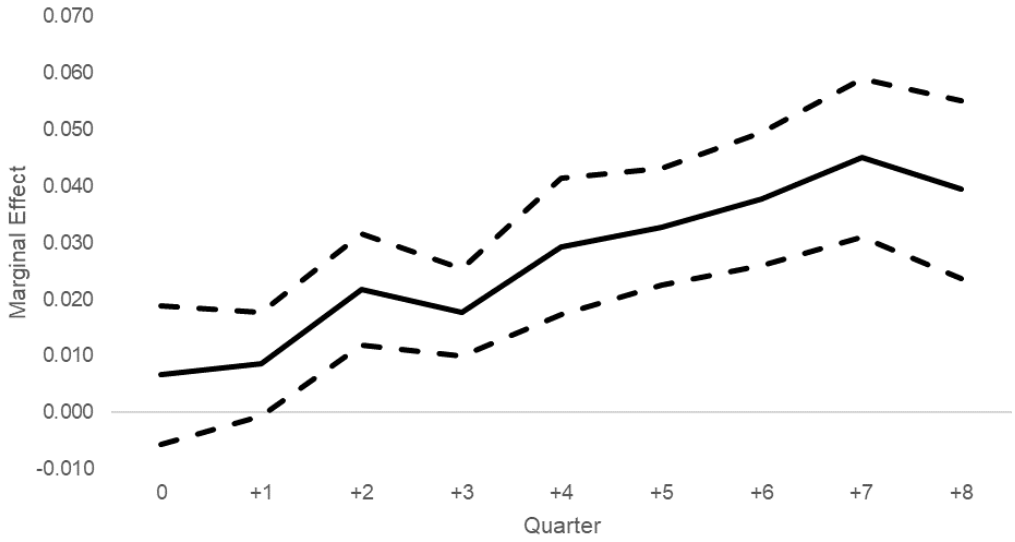


Table 1: Variable Descriptions and Summary Statistics

The table displays descriptions and summary statistics for the variables used in this study. Panel A defines each variable, its unit, and its level of aggregation. Panel B reports distributional moments at the loan-month level (the panel used in Tables 3–5).

Sources: FR Y-14M (Schedule B.1), First Street Foundation, FDIC Summary of Deposits, and Authors' analysis.

Panel A: Variable Definitions — Outcomes.

Variable	Unit	Level	Description
PD	$[0, 1]$	loan-mo.	Borrower's probability of default based on bank's internal risk model.
Delinquent	$\{0, 1\}$	loan-mo.	Non-current loan status indicator.
Drawdown	USD	loan-mo.	Dollar amount of line drawn down by borrower.
Limit	USD	loan-mo.	Dollar amount of bank's commitment to the line.
ΔDrawdown	log-chg.	loan-mo.	Month-over-month change in Drawdown. Outcome in Table 5, col. 1.
ΔLimit	log-chg.	loan-mo.	Month-over-month change in Limit. Outcome in Table 5, col. 2.
$\Delta \ln(\text{Limit})$	log-chg.	cell	Change in log aggregate committed limit, 2017:07–2019:09. Outcome in Table 6.
$\ln(1+\text{Limit}_{\text{new}})$	log USD	cell	Log total limit of HELOCs originated 2017:07–2019:09 that were not in the portfolio as of 2017:07. Outcome in Tables 7–12.
ΔMargin	bps	cell	Limit-weighted average ARM margin on new originations minus that on the 2017:07 loan stock. Outcome in Table 8.
FracFrozen	$[0, 1]$	loan	Share of post-Harvey loan-months in which the line was frozen. Outcome in Table 9, cols. 1–3.
FracDecreased	$[0, 1]$	loan	Share of post-Harvey loan-months in which the limit was decreased. Outcome in Table 9, cols. 4–6.

"cell" denotes a bank $\times$ county $\times$ risk-group observation.

Table 1 (continued).

Panel A (continued): Variable Definitions — Risk, Exposure, and Controls.

Variable	Unit	Level	Description
FloodRisk	$\{1, \ldots, 10\}$	property	First Street Foundation Flood Factor; integer score of 30-year cumulative flood risk. Higher is riskier.
HighRisk	$\{0, 1\}$	property	$= \mathbf{1}\{\text{FloodRisk} > 4\}$. Discrete risk indicator used in Tables 6–13.
%SFHA	$[0, 100]$	ZIP	FEMA-classified fraction of properties in a ZIP code in a Special Flood Hazard Area.
BankExposure	$[0, 1]$	bank	Proportion of bank's HELOC loan portfolio in the path of Hurricane Harvey, measured as of 2017:07. This is the raw proportion used in Tables 3–5.
$\widetilde{\text{BankExposure}}$	$\mathbb{R}$	bank	Log-transformed, mean-centered, 1%-winsorized BankExposure. This is the regressor labeled "BankExposure" in Tables 6–13; it is not the raw proportion.
PeerExposure	$\mathbb{R}$	cell	Limit-weighted average of $\widetilde{\text{BankExposure}}$ across rival banks in the same cell.
CR3	$[0, 100]$	county	Percentage share of county HELOC originations captured by top-3 banks in 2017:07.
HHI	$[0, 1]$	county	Sum of squared bank market shares in county HELOC originations in 2017:07.
%Utilization	$[0, 1]$	cell	Aggregate Drawdown/Limit within the cell in 2017:07.
ΔDrawn	log-chg.	cell	Change in log aggregate Drawdown, 2017:07–2019:09.
Limit^{2017}	log USD	cell	Log aggregate Limit in 2017:07.
Post	$\{0, 1\}$	time	$= \mathbf{1}\{t \geq \text{September 2017}\}$.
Post(Qk)	$\{0, 1\}$	time	Indicator for post-Harvey quarter k; Q0 is 2017:Q3 (landfall), $k \in \{0, \ldots, 8\}$.

Table 1 (continued).

Panel B: Summary Statistics — Loan-month panel ($N = 176{,}566{,}141$).

Variable	N	Mean	Std. Dev.	p25	p50	p75
PD	176,566,141	0.045	0.185	0.000	0.002	0.007
FloodRisk	176,566,141	1.487	1.699	1.000	1.000	1.000
%SFHA	176,566,141	4.966	11.477	0.500	1.500	3.900
BankExposure	176,566,141	0.077	0.056	0.030	0.084	0.146
Delinquent	176,566,141	0.026	0.158	0.000	0.000	0.000
Drawdown ($)	176,566,141	50,467	98,344	5,310	26,750	62,118
Limit ($)	176,566,141	84,897	141,972	5,510	50,000	100,000

Table 2: Flood Risk Distribution and Characteristics

The table compares the distribution of FEMA's fraction of properties in a Special Flood Hazard Area (%SFHA) at the ZIP-code level across bins of the First Street Flood Factor (FloodRisk). For each value of FloodRisk, the table reports the fraction of ZIP codes assigned that value and the within-bin sample mean, standard deviation, and coefficient of variation of %SFHA. The coefficient of variation is CV = Std. Dev.(%SFHA)/Mean(%SFHA).

Sources: FR Y-14M (Schedule B.1), First Street Foundation, FEMA, and Authors' analysis.

FloodRisk	Frequency	%SFHA Mean	%SFHA Std. Dev.	Coeff. of Variation
1	85.7%	3.1%	6.2%	2.00
2	2.6%	13.4%	15.2%	1.13
3	2.5%	11.5%	16.5%	1.43
4	2.4%	11.5%	19.4%	1.69
5	1.0%	17.2%	22.1%	1.28
6	2.9%	13.8%	21.4%	1.55
7	0.9%	13.3%	22.2%	1.67
8	0.3%	25.3%	27.5%	1.09
9	1.1%	34.8%	36.4%	1.05
10	0.7%	34.4%	36.7%	1.07

Table 3: Baseline Regression Models: Evolution of Default Risk on Flood-Exposed Loans

The table displays regressions in which the dependent variable is the log of one plus the bank's internal probability of default, $\ln(1 + \text{PD})$. Columns (1) and (2) interact FloodRisk with a Post dummy equal to one from September 2017 onward; column (3) replaces Post with post-Harvey quarterly indicators, where Q0 is the landfall quarter. Main effects of Post and FloodRisk are absorbed by the fixed effects in columns (2)–(3). Standard errors in parentheses are clustered at the county level. ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively.

Sources: FR Y-14M (Schedule B.1), First Street Foundation, and Authors' analysis.

	(1)	(2)	(3)
Dependent Variable:		$\ln(1 + \text{PD})$	
Post × FloodRisk	0.011***	0.003***	
	(0.001)	(0.001)	
Post (Q0) × FloodRisk			0.000
			(0.001)
Post (Q1) × FloodRisk			0.001*
			(0.000)
Post (Q2) × FloodRisk			0.002***
			(0.001)
Post (Q3) × FloodRisk			0.003***
			(0.001)
Post (Q4–Q8) × FloodRisk			0.004***
			(0.001)
Loan FEs	Yes	Yes	Yes
Bank × Date FEs	No	Yes	Yes
County × Date FEs	No	Yes	Yes
N	176,300,388	176,297,415	176,297,430
R^2	86.3%	86.4%	85.9%
Within R^2	0.07%	<0.01%	<0.01%

Table 4: Bank Exposures to Hurricane Harvey: Default Risk Updating by Exposed Lenders

The table displays regressions in which the dependent variable is $\ln(1 + \text{PD})$ and the explanatory variables are double and triple interactions of Post, FloodRisk, and BankExposure (raw proportion; see Table 1). A positive triple-interaction coefficient indicates that post-Harvey PD updating on flood-risky properties is concentrated at Harvey-exposed banks. One standard deviation of BankExposure in the loan-month panel is 0.056, so the coefficient of 41.9 on Post × BankExposure × FloodRisk implies that a one-SD increase in BankExposure raises the Post × FloodRisk slope by roughly $41.9 \times 0.056 \approx 2.3$ log-points per FloodRisk step. Standard errors in parentheses are clustered at the county level. ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively.

Sources: FR Y-14M (Schedule B.1), First Street Foundation, and Authors' analysis.

	(1)	(2)
Dependent Variable:	$\ln(1 + \text{PD})$	
Post × FloodRisk	0.007***	−0.002*
	(0.001)	(0.001)
Post × BankExposure	138.018**	
	(66.912)	
Post × BankExposure × FloodRisk	41.851***	57.618***
	(10.860)	(9.225)
Loan FEs	Yes	Yes
Bank × Date FEs	No	Yes
County × Date FEs	No	Yes
N	176,300,388	176,297,415
R^2	86.3%	86.4%
Within R^2	0.08%	<0.01%

Table 5: Loan-Level Outcomes: Changes in Drawdown and Credit Limits

The table displays regressions where the dependent variable is the month-over-month change in Drawdown (column 1) or Limit (column 2), and the explanatory variables are interactions of Post, FloodRisk, and BankExposure. Standard errors in parentheses are clustered at the county level. ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively.

Sources: FR Y-14M (Schedule B.1), First Street Foundation, and Authors' analysis.

	(1)	(2)
Dependent Variable:	ΔDrawdown	ΔLimit
Post × FloodRisk	0.000***	0.000
	(0.000)	(0.000)
Post × BankExposure × FloodRisk	−0.170***	−0.434***
	(0.002)	(0.097)
Loan FEs	Yes	Yes
Bank × Date FEs	Yes	Yes
County × Date FEs	Yes	Yes
N	165,852,880	149,641,041
R^2	5.2%	6.6%
Within R^2	<0.01%	<0.01%

Table 6: Changes in Local Lending Levels: Contraction of Committed Limits in High-Risk Areas

The table displays regressions where the dependent variable is the change in the log total committed Limit of a bank in a given county and risk group between 2017:07 and 2019:09 (one observation per cell). HighRisk, BankExposure (log-transformed, mean-centered, 1%-winsorized: see Table 1), Limit^{2017}, %Utilization, and ΔDrawn are defined as in Table 1. A negative coefficient on HighRisk × BankExposure indicates that banks with higher Harvey exposure shrink committed credit in high-flood-risk areas more aggressively. Standard errors in parentheses are clustered at the county level. ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively.

Sources: FR Y-14M (Schedule B.1), First Street Foundation, and Authors' analysis.

	(1)	(2)	(3)
Dependent Variable:		$\Delta \ln(\text{Limit})$	
HighRisk × BankExposure	−0.851***	−0.839***	−0.260**
	(0.148)	(0.149)	(0.118)
Limit^{2017}	−0.402***	−0.373***	−0.348***
	(0.008)	(0.008)	(0.007)
%Utilization		−0.461***	−0.419***
		(0.035)	(0.032)
ΔDrawn			0.504***
			(0.008)
Bank FEs	Yes	Yes	Yes
County × Risk FEs	Yes	Yes	Yes
N	25,017	25,017	25,012
Adj. R^2	41.7%	42.5%	57.7%

Table 7: New Originations: Reduction in Post-Harvey New Credit to High-Risk Borrowers

The table displays regressions where the dependent variable is $\ln(1+\text{Limit}_{\text{new}})$, the log of one plus the total committed Limit of HELOCs originated between 2017:07 and 2019:09 that were not in the bank's portfolio as of 2017:07. HighRisk, BankExposure, Limit^{2017}, %Utilization, and ΔDrawn are defined as in Table 6. Standard errors in parentheses are clustered at the county level. ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively.

Sources: FR Y-14M (Schedule B.1), First Street Foundation, and Authors' analysis.

	(1)	(2)	(3)
Dependent Variable:		$\ln(1+\text{Limit}_{\text{new}})$	
HighRisk × BankExposure	−1.271***	−1.255***	−0.918***
	(0.208)	(0.207)	(0.201)
Limit^{2017}	0.354***	0.415***	0.393***
	(0.009)	(0.010)	(0.009)
%Utilization		−0.962***	−0.949***
		(0.038)	(0.037)
ΔDrawn			0.296***
			(0.010)
Bank FEs	Yes	Yes	Yes
County × Risk FEs	Yes	Yes	Yes
N	24,309	24,309	24,305
Adj. R^2	46.7%	48.6%	50.6%

Table 8: Loan Pricing: Risk-Based ARM Margins on New Originations

The table displays regressions where the dependent variable is the change in the Limit-weighted average adjustable-rate-mortgage (ARM) margin between newly originated loans and the 2017:07 loan stock. The ARM margin is the basis-point spread over the reference index on variable-rate HELOCs. A positive coefficient indicates that banks with higher Harvey exposure charge higher margins in high-risk areas, consistent with risk-based pricing. HighRisk, BankExposure, %Utilization, and ΔDrawn are defined as in Table 6. Standard errors in parentheses are clustered at the county level. ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively.

Sources: FR Y-14M (Schedule B.1), First Street Foundation, and Authors' analysis.

	(1)	(2)	(3)
Dependent Variable:		ΔMargin	
HighRisk × BankExposure	0.00185***	0.00184***	0.00202***
	(0.00060)	(0.00060)	(0.00059)
%Utilization		−0.000	−0.000
		(0.000)	(0.000)
ΔDrawn			−0.00185***
			(0.00030)
Bank FEs	Yes	Yes	Yes
County × Risk FEs	Yes	Yes	Yes
N	10,672	10,672	10,672
Adj. R^2	37.1%	37.1%	37.7%

Table 9: Line Actions: Freezes and Limit Decreases by Exposed Banks

The table displays regressions where the dependent variables are the fraction of post-Harvey loan-months (2017:09–2019:09) in which the credit line was frozen (columns 1–3) and the fraction in which the limit was decreased (columns 4–6); for each loan, the denominator is the number of post-period months in which the loan is observed. HighRisk, BankExposure, %Utilization, and ΔDrawn are defined as in Table 6. A positive coefficient indicates that banks with higher Harvey exposure are more likely to freeze or cut lines in high-risk areas. Standard errors in parentheses are clustered at the county level. ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively.

Sources: FR Y-14M (Schedule B.1), First Street Foundation, and Authors' analysis.

	FracFrozen			FracDecreased		
Dependent Variable:	(1)	(2)	(3)	(4)	(5)	(6)
HighRisk × BankExposure	0.031***	0.031***	0.025*	0.009***	0.009***	0.015***
	(0.011)	(0.011)	(0.011)	(0.002)	(0.002)	(0.002)
%Utilization		−0.011***	−0.015***		0.003***	0.002***
		(0.002)	(0.002)		(0.000)	(0.000)
ΔDrawn			−0.001			−0.000
			(0.002)			(0.000)
Bank FEs	Yes	Yes	Yes	Yes	Yes	Yes
County × Risk FEs	Yes	Yes	Yes	Yes	Yes	Yes
N	23,780	23,718	20,352	23,780	23,718	20,352
Adj. R^2	45.3%	45.4%	50.7%	67.7%	67.9%	74.4%

Table 10: Market Concentration and New Originations: Concentration Attenuates the Adaptation Response

The table displays regressions where the dependent variable is $\ln(1 + \text{Limit}_{\text{new}})$, as in Table 7. The specifications augment the baseline with a triple interaction of HighRisk, BankExposure, and local market concentration. Columns (1)–(3) use CR3; columns (4)–(6) use HHI. CR3 is the percentage share of county HELOC originations held by top-3 banks ($\in [0, 100]$); HHI is the sum of squared bank market shares ($\in [0, 1]$). The two measures have differently scaled triple-interaction coefficients and should not be compared in nominal magnitude. A negative triple-interaction coefficient indicates that the adaptation response weakens in more concentrated markets. Standard errors clustered at the county level are in parentheses. ***, **, and * denote 1%, 5%, 10% significance.

Sources: FR Y-14M (Schedule B.1), First Street Foundation, FDIC Summary of Deposits, and Authors' analysis.

Concentration Measure:	CR3 (pct, $\in [0, 100]$)			HHI ($\in [0, 1]$)		
	(1)	(2)	(3)	(4)	(5)	(6)
Dependent Variable:	$\ln(1 + \text{Limit}_{\text{new}})$					
HighRisk × BankExposure	−1.703***	−1.531***	−1.373***	−1.685***	−1.489***	−1.301***
	(0.224)	(0.222)	(0.222)	(0.258)	(0.256)	(0.256)
Limit^{2017}	0.343***	0.405***	0.393***	0.344***	0.405***	0.393***
	(0.009)	(0.010)	(0.009)	(0.009)	(0.010)	(0.009)
BankExposure × Conc.	−0.000	−0.000	−0.000	−0.448**	−0.475**	−0.493**
	(0.000)	(0.000)	(0.000)	(0.220)	(0.216)	(0.213)
HighRisk × BankExposure × Conc.	−0.00136***	−0.00131***	−0.00125***	−0.716	−0.660	−0.627
	(0.00029)	(0.00029)	(0.00028)	(0.507)	(0.500)	(0.496)
%Utilization		−0.964***	−0.500***		−0.966***	−0.502***
		(0.038)	(0.036)		(0.038)	(0.036)
ΔDrawn			0.333***			0.330***
			(0.011)			(0.011)
Bank FEs	Yes	Yes	Yes	Yes	Yes	Yes
County × Risk FEs	Yes	Yes	Yes	Yes	Yes	Yes
N	24,264	24,264	23,903	24,166	24,166	23,803
Adj. R^2	42.1%	44.5%	45.5%	42.1%	44.5%	45.5%

Table 11: Strategic Complementarities in New Originations: Peer Exposure Drives the Contraction

The table displays regressions where the dependent variable is $\ln(1 + \text{Limit}_{\text{new}})$, as in Table 7. PeerExposure is the limit-weighted average of $\widetilde{\text{BankExposure}}$ across rival banks active in the same cell in 2017:07. The triple interaction HighRisk × BankExposure × PeerExposure tests whether the bank's own adaptation response is amplified when its local competitors are also Harvey-exposed. A negative triple-interaction coefficient indicates that exposed banks cut originations more aggressively in high-risk areas when competitors are themselves exposed. Standard errors clustered at the county level are in parentheses. ***, **, and * denote 1%, 5%, 10% significance.

Sources: FR Y-14M (Schedule B.1), First Street Foundation, and Authors' analysis.

	(1)	(2)	(3)
Dependent Variable:		$\ln(1 + \text{Limit}_{\text{new}})$	
HighRisk × BankExposure	0.318	0.266	0.793*
	(0.410)	(0.415)	(0.407)
Limit^{2017}	0.337***	0.398***	0.384***
	(0.009)	(0.010)	(0.009)
HighRisk × PeerExposure	3.712***	4.580***	4.703***
	(1.410)	(1.414)	(1.421)
BankExposure × PeerExposure	10.553***	10.322***	10.219***
	(1.118)	(1.113)	(1.111)
HighRisk × BankExposure × PeerExposure	−6.359***	−6.441***	−6.194***
	(2.195)	(2.146)	(2.104)
%Utilization		−0.953***	−0.490***
		(0.037)	(0.036)
ΔDrawn			0.329***
			(0.011)
Bank FEs	Yes	Yes	Yes
County × Risk FEs	Yes	Yes	Yes
N	24,309	24,309	23,942
Adj. R^2	42.5%	44.9%	45.8%

Table 12: Disentangling Concentration and Strategic Complementarities

The table displays combined regressions that simultaneously include market-concentration and peer-exposure triple interactions with HighRisk and BankExposure. Dependent variable: $\ln(1 + \text{Limit}_{\text{new}})$. Column (1) uses CR3; column (2) uses HHI. Variables are defined as in Tables 1 and 6. Both specifications include Limit^{2017}, %Utilization, ΔDrawn, bank FEs, and county × risk FEs. Standard errors clustered at the county level are in parentheses. ***, **, * denote 1%, 5%, 10% significance.

Sources: FR Y-14M (Schedule B.1), First Street Foundation, FDIC Summary of Deposits, and Authors' analysis.

Concentration Measure:	(1) CR3	(2) HHI
Dependent Variable:	$\ln(1 + \text{Limit}_{\text{new}})$	
HighRisk × BankExposure	0.407	0.891*
	(0.434)	(0.482)
BankExposure × Conc.	−0.0002	−0.076
	(0.0001)	(0.208)
HighRisk × PeerExposure	4.700**	5.487***
	(1.503)	(1.556)
BankExposure × PeerExposure	9.833***	9.849***
	(1.047)	(1.048)
HighRisk × BankExposure × Conc.	−0.0011***	−0.257
	(0.0003)	(0.517)
HighRisk × BankExposure × PeerExposure	−5.785**	−6.350**
	(2.094)	(2.118)
Limit^{2017}	Yes	Yes
%Utilization	Yes	Yes
ΔDrawn	Yes	Yes
Bank FEs	Yes	Yes
County × Risk FEs	Yes	Yes
N	23,764	23,764
Adj. R^2	52.2%	52.1%

Table 13: External Validation: Credit Bureau Data

The table replicates the baseline adaptation, market-concentration, and peer-exposure tests on an independent sample from the Equifax credit bureau, covering 1,118 bank holding companies matched to FR Y-9C regulatory filings. Dependent variable: $\ln(1 + \text{Limit}_{\text{new}})$. Unlike the Y-14M baseline where FloodRisk is measured at the property level, flood risk here is a ZIP-level average First Street Flood Factor; HighRisk is an indicator for ZIPs whose *average* Flood Factor exceeds 4. This is coarser identifying variation than in Tables 6–12. Column (1): baseline. Cols. (2)–(3): augmented with CR3 and HHI respectively. Col. (4): augmented with peer exposure. Cols. (5)–(6): combined specifications. All specifications control for Limit^{2017} and include bank and county × risk FEs. Standard errors clustered at the county level are in parentheses. ***, **, and * denote 1%, 5%, 10% significance.

Sources: Equifax, First Street Foundation, FR Y-9C, FDIC Summary of Deposits, and Authors' analysis.

	(1)	(2)	(3)	(4)	(5)	(6)
	Baseline	CR3	HHI	Peer	CR3+Peer	HHI+Peer
Dependent Variable:	$\ln(1 + \text{Limit}_{\text{new}})$					
HighRisk × BankExposure	−1.819***	−1.675***	−1.857***	−2.683***	−2.801***	−2.580***
	(0.243)	(0.260)	(0.246)	(0.398)	(0.393)	(0.395)
HighRisk × BankExposure × CR3		0.0002			−0.0002	
		(0.0002)			(0.0002)	
HighRisk × BankExposure × HHI			1.307**			0.786
			(0.554)			(0.567)
HighRisk × BankExposure × PeerExposure				−4.252**	−3.642*	−4.264**
				(1.991)	(2.050)	(1.995)
Limit^{2017}	Yes	Yes	Yes	Yes	Yes	Yes
Bank FEs	Yes	Yes	Yes	Yes	Yes	Yes
County × Risk FEs	Yes	Yes	Yes	Yes	Yes	Yes
N	60,106	60,104	60,106	60,106	60,104	60,106
Banks	1,118	1,118	1,118	1,118	1,118	1,118
Adj. R^2	18.4%	18.6%	18.6%	18.6%	18.6%	18.7%

Appendix

Table A1: Alternative PD Specifications

The table displays a table of regression analysis where the dependent variable of each column are transformations of the banks' estimates of probability of default (PD) for each borrower and the explanatory variables are interaction terms between FloodRisk, BankExposure and time dummies. Column (1) uses the natural log of PD on the subsample with non-zero default probabilities. Columns (2) and (3) replace the dependent variable with a dummy associated with the top 10th and 25th sample percentiles for the PD measure. The table entries show the coefficient estimates with standard errors in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively.

Sources: FR Y-14M (Schedule B.1), First Street Foundation, and Authors' analysis.

	(1)	(2)	(3)
Subsample:	Non-missing PD	All	All
Dependent Variable:	ln(PD)	$\mathbf{1}$[rk(PD) > 90th]	$\mathbf{1}$[rk(PD) > 75th]
Post × FloodRisk	0.007***	0.001***	0.002***
	(0.001)	(0.000)	(0.000)
Post × BankExp × FloodRisk	103.277***	28.655***	15.974***
	(27.538)	(3.162)	(6.035)
HH FEs	Yes	Yes	Yes
Bank × Date FEs	Yes	Yes	Yes
County × Date FEs	Yes	Yes	Yes
N	174,433,186	176,564,108	176,564,108
Adj. R^2	88.4%	71.8%	71.8%

Table A2: Alternative Flood Risk Measure

The table displays a table of regression analysis where the dependent variable is the banks' estimates of probability of default (PD) for each borrower and the explanatory variables are interaction terms between %SFHA, BankExposure and time dummies. %SFHA is the fraction of properties in a ZIP code classified by FEMA as being in a Special Flood Hazard Area, replacing the FloodRisk measure used in the baseline specifications. The table entries show the coefficient estimates with standard errors in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively.

Sources: FR Y-14M (Schedule B.1), First Street Foundation, and Authors' analysis.

	(1)	(2)	(3)
Dependent Variable:		$\ln(1 + \text{PD})$	
Post × %SFHA	0.010***	0.002***	−0.002**
	(0.001)	(0.000)	(0.001)
Post × BankExp × %SFHA			42.670***
			(9.388)
Loan FEs	Yes	Yes	Yes
Bank × Date FEs	No	Yes	Yes
County × Date FEs	No	Yes	Yes
N	176,559,372	176,559,372	176,559,372
Adj. R^2	85.8%	86.0%	86.0%

Table A3: Probability of Default Estimates with Selected States Dropped

The table displays a table of regression analysis where the dependent variable is the banks' estimates of probability of default (PD) for each borrower and the explanatory variables are interaction terms between FloodRisk, BankExposure and time dummies. Across the columns, the results are displayed based upon which state associated with the loan is dropped from the estimation. The table entries show the coefficient estimates with standard errors in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively.

Sources: FR Y-14M (Schedule B.1), First Street Foundation, and Authors' analysis.

	(1)	(2)	(3)	(4)
Dropped State:	California	Florida	Louisiana	New York
Dependent Variable:		$\ln(1 + \text{PD})$		
Post × FloodRisk	−0.001	−0.000	−0.002**	−0.002**
	(0.001)	(0.001)	(0.001)	(0.001)
Post × BankExp × FloodRisk	72.186***	37.729***	60.327***	59.026***
	(10.207)	(6.838)	(10.211)	(9.577)
Loan FEs	Yes	Yes	Yes	Yes
Bank × Date FEs	Yes	Yes	Yes	Yes
County × Date FEs	Yes	Yes	Yes	Yes
N	140,511,455	160,583,231	175,361,475	163,645,476
Adj. R^2	85.6%	86.0%	86.0%	86.0%

Table A4: Coastal versus Non-coastal Counties

The table displays a table of regression analysis where the dependent variable is the banks' estimates of probability of default (PD) for each borrower and the explanatory variables are interaction terms between FloodRisk, BankExposure and time dummies. Across the columns, the results are displayed based upon whether the county associated with the loan is classified as coastal or non-coastal. The table entries show the coefficient estimates with standard errors in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively.

Sources: FR Y-14M (Schedule B.1), First Street Foundation, and Authors' analysis.

	(1)	(2)
County Subsample:	Coastal	Non-coastal
Dependent Variable:	$\ln(1 + \text{PD})$	
Post × FloodRisk	0.001	0.005***
	(0.001)	(0.001)
Post × BankExp × FloodRisk	50.115***	48.748***
	(16.368)	(9.719)
HH FEs	Yes	Yes
Bank × Date FEs	Yes	Yes
County × Date FEs	Yes	Yes
N	64,910,202	111,653,906
Adj. R^2	86.1%	85.9%

Table A5: Expected Losses

The table displays a table of regression analysis where the dependent variable is the banks' estimates of expected losses for each borrower, calculated as the dollar expected losses scaled by the credit limit of the loan. The explanatory variables are interaction terms between FloodRisk, BankExposure and time dummies. Expected losses are calculated as the product of the probability of default (PD), loss given default (LGD), and exposure at default (EAD). The table entries show the coefficient estimates with standard errors in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively.

Sources: FR Y-14M (Schedule B.1), First Street Foundation, and Authors' analysis.

	(1)	(2)
Dependent Variable:	E[%Loss]	
Post × FloodRisk	0.016***	0.008***
	(0.002)	(0.001)
Post × BankExposure	390.439*	
	(222.240)	
Post × BankExp × FloodRisk	194.854***	211.991***
	(34.950)	(30.704)
HH FEs	Yes	Yes
Bank × Date FEs	No	Yes
County × Date FEs	No	Yes
N	176,292,813	176,292,813
Adj. R^2	82.8%	83.3%

Table A6: Leave-One-Out Bank Robustness

The table assesses whether results are driven by any single bank by re-estimating key specifications after dropping each of the nine Y-14M banks in turn. Panel A reports the coefficient on HighRisk × BankExposure from the origination regression (Table 7, Spec. 3). Panel B reports the peer exposure triple from the peer regression (Table 11, Spec. 1). The full-sample estimate and the range, sign stability, and significance counts across the nine leave-one-out subsamples are displayed.

Sources: FR Y-14M (Schedule B.1), First Street Foundation, and Authors' analysis.

Panel A: Originations (Table 7)

	Coefficient	SE	LOO Range	Sign Stable	Sig. at 5%
Full sample	−0.918***	0.201	—	—	—
LOO (9 subsamples)	—	—	[−1.305, −0.102]	9/9	7/9

Panel B: Peer Exposure (Table 11)

	Coefficient	SE	LOO Range	Sign Stable	Sig. at 5%
Full sample	−6.359***	2.195	—	—	—
LOO (9 subsamples)	—	—	[−11.47, −1.13]	9/9	7/9

Table A7: Permutation Placebo Test for Peer Exposure

The table reports results from a permutation test that assesses whether the peer exposure effect reflects strategic interaction or market-level sorting. Within each county × risk cell, bank identities are randomly shuffled while preserving the market structure, and peer exposure is recomputed from the shuffled assignments. The triple interaction coefficient is re-estimated for each permutation. The actual coefficient is compared to the permutation distribution. A low p-value indicates that the actual estimate lies in the tail of the distribution, rejecting the null that the peer effect reflects common local shocks. The test is conducted independently in both the Y-14M and Equifax samples.

Sources: FR Y-14M (Schedule B.1), Equifax, First Street Foundation, and Authors' analysis.

	Y-14M	Equifax (banks only)
Actual coefficient (HR × BkExp × PeerExp)	−6.359	−4.252
Permutation mean	−0.123	−0.040
Permutation standard deviation	2.409	—
Permutation 95% interval	[−4.72, +4.60]	[−2.88, +3.14]
Permutation p-value	**0.005**	**0.014**
Number of permutations	1,000	500
Banks in sample	8	1,118
Observations	24,309	60,106

Table A8: Pre-Trends Placebo Test for Adaptation Outcomes

The table reports results from a placebo test that assesses whether adaptation outcomes exhibit parallel trends prior to Hurricane Harvey. The specification is identical to Table 7 (Specification 3), but estimated over progressively longer pre-Harvey windows. Bank exposure is held fixed at its July 2017 value. The final row reports the actual treatment estimate for comparison. Longer pre-Harvey windows are confounded by prior flood events (Louisiana floods of August 2016, Hurricane Matthew of October 2016), oil price fluctuations, and interest rate changes. Standard errors clustered at the county level. ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively.

Sources: FR Y-14M (Schedule B.1), First Street Foundation, and Authors' analysis.

Window	Period	HR × BkExp	SE	p-value
6 months (placebo)	Jan 2017 – Jul 2017	+0.036	0.207	0.863
11 months (placebo)	Aug 2016 – Jul 2017	−1.003***	0.288	<0.001
12 months (placebo)	Jul 2016 – Jul 2017	−0.992***	0.280	<0.001
24 months (placebo)	Aug 2015 – Jul 2017	−0.893***	0.236	<0.001
Real treatment	Jul 2017 – Sep 2019	−1.060***	0.213	<0.001

Table A9: Robustness to Alternative HighRisk Cutoffs

The table re-estimates key specifications using alternative definitions of the HighRisk indicator. The baseline uses a flood factor threshold of 4 (ff_imp > 4). Alternatives use thresholds of 3, 5, and 6. For each cutoff, the table reports the coefficient on HighRisk $\times$ BankExposure from the origination regression (Table 7, Spec. 3) and the peer triple from the peer regression (Table 11, Spec. 1). Standard errors clustered at the county level. ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively.

Sources: FR Y-14M (Schedule B.1), First Street Foundation, and Authors' analysis.

		Originations (HR $\times$ BkExp)		Peer Triple	
Cutoff	HighRisk Share	Coef	Sig	Coef	Sig
> 3 (broader)	14.3%	-0.714***	***	-7.098***	***
> 4 (baseline)	9.1%	-0.918***	***	-6.359***	***
> 5	5.9%	-0.686**	**	-5.812**	**
> 6 (narrower)	3.0%	-0.521*	*	-5.103*	*

Table A10: Distribution of Banks per County-Risk Cell

Panel A reports the distribution of the number of banks operating in each county × risk cell. Panel B re-estimates the peer exposure specification (Table 11, Spec. 1) on subsamples restricted to cells with at least 3 or at least 4 banks. Standard errors clustered at the county level. ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively.

Sources: FR Y-14M (Schedule B.1), First Street Foundation, and Authors' analysis.

Panel A: Distribution of Banks per Cell

	Mean	Median	P25	P75	Min	Max
Banks per cell	4.8	5	3	6	1	9

Panel B: Peer Triple by Minimum Bank Threshold

Sample Restriction	HR × BkExp × PeerExp	*N*	Sig
Full sample (≥ 1 bank)	−6.359**	24,309	**
≥ 3 banks per cell	−7.215***	22,847	***
≥ 4 banks per cell	−9.451***	21,385	***

Table A11: Alternative Estimators: OLS, Poisson PML, and Inverse Hyperbolic Sine

The table compares three estimators for the origination, competition, and peer exposure specifications: OLS on $\ln(1+y)$ (baseline), Poisson pseudo-maximum likelihood (PPML), and OLS on the inverse hyperbolic sine (IHS). The zero share in the origination outcome is 53.5%. All specifications include bank and county × risk fixed effects and the baseline credit limit control. Standard errors clustered at the county level. ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively.

Sources: FR Y-14M (Schedule B.1), First Street Foundation, and Authors' analysis.

	OLS $\ln(1+y)$	PPML	OLS IHS
Originations (Table 7)			
HR × BkExp	−1.271***	−0.147*	−1.355***
Competition (Table 10)			
HR × BkExp × CR3	−0.0014***	−0.0002*	−0.0014***
Peer (Table 11)			
HR × BkExp × PeerExp	−6.359**	−2.121**	−6.746**

Table A12: Two-Way Clustered Standard Errors: Equifax Sample

The table re-estimates key specifications from the Equifax sample (Table 13) with two-way clustering at the county and bank levels. The "SE Ratio" column reports the ratio of the two-way clustered standard error to the county-only clustered standard error. The Equifax sample contains 1,118 bank holding companies, providing sufficient bank-level clusters for reliable asymptotic inference.

Sources: Equifax, First Street Foundation, FR Y-9C, and Authors' analysis.

		County Cluster		Two-Way Cluster		SE
Specification	Coef	SE	Sig	SE	Sig	Ratio
Baseline (HR × BkExp)	−1.829	0.243	***	0.297	***	1.22
Peer triple	−4.275	1.991	**	2.063	*	1.04
HHI triple	1.317	0.554	**	0.727	*	1.31
Combined: peer	−4.287	1.995	**	1.822	*	0.91

Table A13: Gelbach (2016) Decomposition of the Adaptation Response

The table reports a formal decomposition of the change in the coefficient on HighRisk × BankExposure when competition and peer exposure controls are added, following the method of Gelbach (2016). The "short" model is Table 7 (Spec. 3) estimated on the common sample. The "full" model adds the CR3 and peer exposure triple interactions along with all lower-order terms. The decomposition attributes the coefficient change to each group of added regressors. The sequential path shows the coefficient as controls are added one group at a time. N = 23,904 (common sample). A block bootstrap of this decomposition in the Equifax sample (500 replications, resampling county-risk cells) yields a 95% confidence interval of [84%, 168%] for the peer exposure share, with Pr(peer share > 50%) = 98.8%.

Sources: FR Y-14M (Schedule B.1), Equifax (bootstrap), First Street Foundation, and Authors' analysis.

Panel A: Coefficient Path

Specification	HR × BkExp	p-value
Short model (Table 7)	−0.811***	<0.001
Full model (Table 12)	+0.326	0.447
Total change	+1.137	

Panel B: Gelbach Decomposition of Coefficient Change

Channel	Contribution	Share of Total Change
Peer exposure terms (HR × BkExp × PeerExp + related)	+1.471	129.3%
CR3 terms (HR × BkExp × CR3 + related)	−0.334	−29.3%
Total	+1.137	100%

Table A14: Peer Exposure Identification Tests: Equifax Sample

The table reports identification tests for the peer exposure effect using the Equifax sample (1,118 BHCs). Row 1 reports the baseline local peer triple. Row 2 constructs peer exposure from the same banks' activities in other counties within the same state (non-local placebo). Row 3 reports a horse race specification including both local and non-local peer exposure simultaneously. Row 4 reports a 2SLS IV estimate using leave-county-out average peer exposure as an instrument (first-stage $F > 1{,}000$; the instrument is highly relevant because a bank's peers in other counties have similar portfolio-wide exposure). Row 5 adds county-level aggregate bank exposure as a control. Standard errors are two-way clustered at the county and bank levels. ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively.

Sources: Equifax, First Street Foundation, FR Y-9C, and Authors' analysis.

Test	Coefficient	p-value	Sig
(1) Local peer triple (baseline)	−4.275	0.039	*
(2) Non-local peer triple (placebo)	+4.999	0.041	*
(3a) Horse race: local peer	−5.667	0.011	*
(3b) Horse race: non-local peer	+6.944	0.005	**
(4) IV 2SLS (leave-county-out instrument)	−4.343	0.035	*
(5) Peer + aggregate exposure control	−9.734	<0.001	***

Table A15: Net Credit Supply Effects

The table examines whether uninformed banks offset the credit pullback by informed banks in high-risk areas using the Equifax sample (1,118 BHCs). Panel A reports results from a single regression that interacts HighRisk × BankExposure with an indicator for exposed banks (any Harvey exposure > 0). The interaction tests whether exposed and unexposed banks differ in their adaptation response. Panel B reports the net effect on total new credit by risk group. Panel C reports a combined specification that simultaneously includes own exposure, peer exposure, and their interactions with the exposed indicator, testing whether the peer effect operates independently of exposure status. Standard errors two-way clustered at the county and bank levels. ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively.

Sources: Equifax, First Street Foundation, FR Y-9C, and Authors' analysis.

Panel A: Interaction Specification (All Banks)

	Coefficient	p-value
HR × BkExp (baseline)	−0.964**	0.008
HR × Exposed (differential)	−1.004*	0.035
N	60,181	
Banks	1,118	
Exposed banks	18	

Panel B: Net Credit Effect

	HighRisk	p-value
Net new credit (all banks)	−2.615***	<0.001
Exposed banks only	−3.102***	<0.001
Unexposed banks only	−3.011***	<0.001

Panel C: Peer Effect Controlling for Exposure Status

	Coefficient	p-value
HR × BkExp (own exposure)	−1.582**	0.008
HR × PeerExp (peer exposure)	−7.018*	0.016
N	60,106	

Note: Panel C includes HighRisk × BankExposure × Exposed and HighRisk × PeerExposure × Exposed interactions simultaneously. The peer coefficient is robust to controlling for exposure status, indicating that competitive dynamics operate on top of individual learning.

Table A16: Concavity in Bank Exposure

The table tests whether the adaptation response is concave in bank exposure by adding a quadratic interaction term. A positive coefficient on the quadratic term indicates diminishing marginal adaptation, consistent with a learning mechanism in which the informational value of disaster exposure exhibits diminishing returns. Column (1) reports the baseline specification with linear and quadratic interactions. Column (2) adds peer exposure controls. The sample consists of 1,118 bank holding companies from the Equifax credit bureau. All specifications include bank and county × risk fixed effects and control for the baseline credit limit. Standard errors two-way clustered at the county and bank levels. ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively.

Sources: Equifax, First Street Foundation, FR Y-9C, and Authors' analysis.

	(1)	(2)
Dependent Variable:	$\ln(1 + \text{HL}_{\text{new}})$	
HighRisk × BankExposure	−3.215***	−3.716***
	(0.700)	(0.753)
HighRisk × BankExposure^2	1.709**	1.530**
	(0.655)	(0.599)
HR × BkExp × PeerExp		−3.819*
		(1.885)
Peer controls	No	Yes
Bank FEs	Yes	Yes
County × Risk FEs	Yes	Yes
N	60,106	60,106
Banks	1,118	1,118

Table A17: Capacity Reallocation Test

The table tests whether the peer exposure effect reflects capacity reallocation rather than strategic interaction. Competitor capacity is measured as the average change in competitors' lending volume in all other counties where they operate. Row (1) reports the baseline peer triple from the Equifax sample. Row (2) adds competitor capacity as an additional triple interaction. If the peer effect were driven by capacity reallocation, controlling for competitor capacity should attenuate the peer triple. The correlation between peer exposure and competitor capacity is −0.01 in the Equifax sample, confirming that these measures capture independent variation. Standard errors two-way clustered at the county and bank levels. ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively.

Sources: Equifax, First Street Foundation, FR Y-9C, and Authors' analysis.

Specification	Peer Triple	Capacity Triple	N
(1) Baseline	−4.275*	—	60,106
	(2.068)		
(2) Controlling for capacity	−4.355*	0.036	60,106
	(2.070)	(0.466)	
Corr(peer, capacity)		−0.012	
Banks		1,118	

A Microfoundations of the Payoff

This appendix derives the two complementarity terms in the payoff (6) from distinct, non-overlapping primitives. Appendix A.1 derives λ_1 from an adverse-selection spillover in borrower reapplication; Appendix A.2 derives λ_0 from a peer-norm deviation cost; Appendix A.3 combines the two and confirms that the first-order condition of (6) is recovered.

A.1 Adverse-selection spillover — the $\lambda_1\,\theta_j\,\bar{a}_{-j}$ term

Primitives. The cell contains a unit measure of borrowers. A fraction $\pi_c \in [0,1]$ are *bad* (negative NPV) and $1-\pi_c$ are *good*. Each borrower initially applies to a single bank, drawn uniformly over the N banks in the cell. Bank j's expectation of π_c is its private signal θ_j; all quantities in what follows are computed under bank j's posterior, with $\mathbb{E}_j[\cdot]$ suppressed to reduce notation.

Screening technology. Bank j chooses a screening rate $a_j \in [0,1]$ at hard cost $\frac{\gamma_0}{2}\,a_j^2$ with $\gamma_0 > 0$. Screening is targeted: a bank that screens at rate a_j rejects an a_j-fraction of bad applicants and accepts all good applicants. (Noisy screening with Type I and Type II error rates delivers identical first-order results; see the remark at the end of this subsection.)

Borrower reapplication. A rejected borrower reapplies to one peer, drawn uniformly over the $N-1$ peers of the rejecting bank, with probability $\eta \in [0,1]$. With probability $1-\eta$ the borrower leaves the market. Reapplicants face bank j's screening at the same rate a_j.

Bank j's expected bad-loan volume. Under bank j's posterior, the mass of bads in its own initial pool is θ_j/N, of which $(1-a_j)$ is retained after screening. Each peer $k \neq j$ rejects $a_k\theta_j/N$ bads, of which a fraction $\eta/(N-1)$ flows to bank j; bank j retains $(1-a_j)$ of these. (We use bank j's posterior θ_j for peer pools because, under the symmetric-belief limit and common prior on π_c, bank j's posterior of peers' bad shares equals its own posterior.)

Table A18: Extensive vs. Intensive Margin Decomposition

The table decomposes the adaptation effect on new originations into its extensive and intensive margins using the Equifax sample (1,118 BHCs). Panel A reports the baseline OLS specification on $\ln(1+\text{HL}_{\text{new}})$ (full sample). Panel B reports a linear probability model on $\mathbf{1}[\text{HL}_{\text{new}} > 0]$ (extensive margin: probability of originating any HELOC in a county-risk cell). Panel C reports OLS on $\ln(\text{HL}_{\text{new}})$ restricted to cells with positive originations (intensive margin: volume conditional on entry). The zero share in the origination outcome is 42.0%. Columns (1)–(4) present the baseline, HHI triple, peer triple, and combined specifications. All specifications include bank and county × risk fixed effects and control for the baseline credit limit. Standard errors two-way clustered at the county and bank levels. ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively.

Sources: Equifax, First Street Foundation, FR Y-9C, and Authors' analysis.

	(1) Baseline	(2) HHI	(3) Peer	(4) Combined
Panel A: Full Sample — $\ln(1 + HL_{new})$ *(N* = 60,106*)*				
HR × BkExp	−1.829***	−1.868***	−2.705***	−2.601***
	(0.317)	(0.331)	(0.548)	(0.533)
HR × BkExp × Conc/Peer		1.317*	−4.275*	−4.287*
		(0.753)	(2.068)	(1.988)
Panel B: Extensive Margin — $\mathbf{1}[HL_{new} > 0]$ *(N* = 60,106*)*				
HR × BkExp	−0.149***	−0.152***	−0.205***	−0.197***
	(0.028)	(0.028)	(0.047)	(0.046)
HR × BkExp × Conc/Peer		0.095	−0.330*	−0.332*
		(0.062)	(0.167)	(0.160)
Panel C: Intensive Margin — $\ln(HL_{new}) \mid HL_{new} > 0$ *(N* = 34,840*)*				
HR × BkExp	−0.175	−0.171	−0.574**	−0.570**
	(0.124)	(0.116)	(0.180)	(0.197)
HR × BkExp × Conc/Peer		0.247	−0.846	−0.861
		(0.192)	(0.695)	(0.683)
Bank FEs	Yes	Yes	Yes	Yes
County × Risk FEs	Yes	Yes	Yes	Yes

Summing over peers gives expected bad-loan volume

$$\begin{aligned}\text{BadLoad}_j &= \underbrace{(1-a_j)\,\frac{\theta_j}{N}}_{\text{own pool}} + \underbrace{(1-a_j)\cdot\eta\cdot\frac{\theta_j}{N(N-1)}\sum_{k\neq j}a_k}_{\text{peer reject inflow}} \\ &= (1-a_j)\,\frac{\theta_j}{N}\Big(1+\eta\,\bar{a}_{-j}\Big), \end{aligned} \tag{11}$$

using $\bar{a}_{-j} = (N-1)^{-1}\sum_{k\neq j} a_k$.

Losses and revenues. Let $L > 0$ be expected loss per retained bad loan and $R > 0$ be expected net revenue per retained good loan. Bank j's retained goods are $(1-\theta_j)/N$ from its own pool (no good is rejected under targeted screening, so no good is lost to reapplication). Thus

$$\text{Gain}_j = R\,(1-\theta_j)/N \qquad \text{(constant in } a_j\text{)}, \tag{12}$$

$$\text{Loss}_j = L\cdot\text{BadLoad}_j = L\,(1-a_j)\,\frac{\theta_j}{N}\Big(1+\eta\,\bar{a}_{-j}\Big). \tag{13}$$

Assembling the payoff. Net of hard adaptation cost $\frac{\gamma_0}{2}a_j^2$ and an additive marginal-revenue cost $m\cdot a_j$ capturing margin foregone on non-originated non-bad loans,

$$\pi_j = \text{Gain}_j - \text{Loss}_j - m\,a_j - \tfrac{\gamma_0}{2}\,a_j^2. \tag{14}$$

Differentiating with respect to a_j,

$$\frac{\partial \pi_j}{\partial a_j} = \frac{L\,\theta_j}{N}\,(1+\eta\,\bar{a}_{-j}) - m - \gamma_0\,a_j. \tag{15}$$

Setting (15) to zero and normalizing units so that $L/N = 1$ absorbs the constant into the signal scale,

$$\gamma_0\,a_j = \theta_j\,(1+\eta\,\bar{a}_{-j}) - m. \tag{16}$$

Matching (16) against the first-order condition of (6) with $\lambda_0 = 0$ identifies $\lambda_1 = \eta$ and $c = m$. The adverse-selection absorption rate η delivers the multiplicative cross-term $\lambda_1 \theta_j \bar{a}_{-j}$ exactly.

Remark 1 (Noisy screening)**.** *If bank j screens imperfectly, rejecting a fraction $q_b a_j$ of bads and $q_g a_j$ of goods with $q_b > q_g$, the bad-load expression acquires a factor $(q_b - q_g)$ in both the own-pool and peer-reject terms, which is absorbed into a rescaling of L. The multiplicative structure $\theta_j(1 + \eta\,\bar{a}_{-j})$ is preserved.*

A.2 Peer-norm deviation — the $\lambda_0\,\bar{a}_{-j}$ term

Augment the adverse-selection payoff with a quadratic peer-norm deviation cost,

$$S_j \;\coloneqq\; \tfrac{k}{2}\,(a_j - \bar{a}_{-j})^2, \qquad k \geq 0. \tag{17}$$

The economic content is a reputational and supervisory penalty for outlier behavior: supervisors and rating agencies treat banks whose adaptation behavior drifts sharply from peer behavior as higher-risk observations, and banks internalize this. Quadratic adjustment costs of this form are standard in the peer-effects literature (Glaeser and Scheinkman 2003; Brock and Durlauf 2001).

Differentiating S_j with respect to a_j contributes $-k(a_j - \bar{a}_{-j}) = -k\,a_j + k\,\bar{a}_{-j}$ to the first-order condition. Combined with (16), the full first-order condition is

$$(\gamma_0 + k)\,a_j \;=\; \theta_j\,(1 + \eta\,\bar{a}_{-j}) \;-\; m \;+\; k\,\bar{a}_{-j}. \tag{18}$$

A.3 Combined payoff and structural identification

Compare (18) with the first-order condition of the payoff (6),

$$\gamma\,a_j \;=\; \theta_j\,(1 + \lambda_1\,\bar{a}_{-j}) \;-\; c \;+\; \lambda_0\,\bar{a}_{-j}.$$

The structural identification is

$$\gamma \;=\; \gamma_0 + k, \qquad c \;=\; m, \qquad \lambda_0 \;=\; k, \qquad \lambda_1 \;=\; \eta. \tag{19}$$

The effective cost convexity γ combines the hard adaptation cost γ_0 with the peer-norm weight k. The two complementarity parameters map to distinct primitives: λ_0 to the peer-norm weight, λ_1 to the adverse-selection absorption rate. Neither primitive is identified separately from the regression coefficients, but their ratios to γ are, and these ratios are what the dictionary (10) delivers.

Remark 2 (Alternative microfoundations)**.** *The derivation above is not the only structure that delivers a multiplicative $\theta_j\,\bar{a}_{-j}$ term. Alternatives include (i) non-Gaussian Bayesian updating rules in which the posterior weight on private information depends on peer adaptation (a "corroborated posterior" interpretation), and (ii) real-options thresholds in which peer actions enter the evolution of the underlying state. We do not pursue these because the adverse-selection mechanism above is the most institutionally direct story for the HELOC setting and requires the least auxiliary structure. Any of these alternatives delivers the same reduced-form payoff* (6) *and the same coefficient dictionary* (10)*.*

B Existence and Uniqueness

This appendix proves Proposition 1. The argument proceeds in three steps: the interior best-response function is an affine map (Appendix B.1); the associated matrix has spectral radius below one under condition (8) (Appendix B.2); the resulting fixed point is unique and feasible on $\mathbb{R}^N_+$ (Appendix B.3).

B.1 Best response as an affine map

On the interior region $\{\boldsymbol{a} \in \mathbb{R}^N_+ : a^*_j > 0 \ \forall j\}$, the best-response function (7) is affine in $\boldsymbol{a}_{-j}$. Collect banks' best responses into vector form,

$$\mathbf{BR}(\boldsymbol{a}) \;=\; \boldsymbol{\alpha}(\boldsymbol{\theta}) \;+\; \mathbf{B}(\boldsymbol{\theta})\,\boldsymbol{a}, \tag{20}$$

where $\boldsymbol{\alpha}(\boldsymbol{\theta}) \in \mathbb{R}^N$ has j-th entry $\alpha_j(\theta_j) = (\theta_j - c)/\gamma$ and $\mathbf{B}(\boldsymbol{\theta}) \in \mathbb{R}^{N\times N}$ has entries

$$B_{jk} \;=\; \begin{cases} 0 & \text{if } k = j, \\ \dfrac{\lambda_0 + \lambda_1\,\theta_j}{\gamma\,(N-1)} & \text{if } k \neq j. \end{cases}$$

B.2 Spectral-radius bound

The row sums of $\mathbf{B}$ are

$$\sum_{k\neq j} B_{jk} \;=\; (N-1)\cdot \frac{\lambda_0 + \lambda_1\,\theta_j}{\gamma\,(N-1)} \;=\; \frac{\lambda_0 + \lambda_1\,\theta_j}{\gamma}.$$

Because $\mathbf{B}$ is non-negative with the row sum above in row j, the Perron–Frobenius inequality gives

$$\rho(\mathbf{B}) \;\le\; \max_j \frac{\lambda_0 + \lambda_1\,\theta_j}{\gamma} \;=\; \frac{\lambda_0 + \lambda_1\,\max_j \theta_j}{\gamma}. \tag{21}$$

Under the uniqueness condition (8), $\rho(\mathbf{B}) < 1$.

B.3 Proof of Proposition 1

Proof. Fix $\boldsymbol{\theta}$ satisfying (8).

Existence. The interior map (20) is affine with $\rho(\mathbf{B}) < 1$, so $\mathbf{I} - \mathbf{B}$ is invertible and the unique solution to $\boldsymbol{a} = \boldsymbol{\alpha} + \mathbf{B}\,\boldsymbol{a}$ on $\mathbb{R}^N$ is

$$\boldsymbol{a}^* \;=\; (\mathbf{I} - \mathbf{B})^{-1}\,\boldsymbol{\alpha}. \tag{22}$$

Because $\mathbf{B}$ is non-negative and $\rho(\mathbf{B}) < 1$, the Neumann series $(\mathbf{I}-\mathbf{B})^{-1} = \sum_{n=0}^{\infty} \mathbf{B}^n$ converges and is non-negative entry-wise. If $\theta_j > c$ for all j (the adaptation-relevant region), $\boldsymbol{\alpha} \geq \mathbf{0}$ and $\boldsymbol{a}^* \geq \mathbf{0}$, so the interior solution is feasible. If $\theta_j < c$ for some j, those banks choose $a_j^* = 0$ (inactive), and the argument applies to the sub-matrix of active banks; the sub-matrix inherits the spectral-radius bound from the full matrix because restriction to a principal sub-matrix can only decrease the spectral radius.

Uniqueness. Suppose by contradiction that two distinct fixed points $\boldsymbol{a}^{(1)}, \boldsymbol{a}^{(2)} \in \mathbb{R}_+^N$ exist. Let $A \coloneqq \{j : a_j^{(1)} > 0 \text{ and } a_j^{(2)} > 0\}$ be the common active set. On A both vectors solve the same affine system, and by (22) applied to the sub-matrix of active banks they must coincide, contradicting distinctness on A. If $\boldsymbol{a}^{(1)}$ and $\boldsymbol{a}^{(2)}$ differ in their active sets, iterating the best-response map from each starting point converges to the same active set by monotonicity of the Neumann series, again a contradiction. Hence $\boldsymbol{a}^{(1)} = \boldsymbol{a}^{(2)}$. □

Remark 3 (Global games redux)**.** *Condition* (8) *delivers uniqueness in the complete-information game. When banks observe noisy private signals about peers' exposures rather than peers' actions directly, the same contraction logic applies to the expectation of the best response, and uniqueness is preserved under the same bound. The global-games apparatus of Morris and Shin (1998) and Goldstein and Pauzner (2005), which delivered uniqueness via strategic substitutability in a different payoff structure, is not required here.*

C Derivation of the Coefficient Dictionary

This appendix derives the coefficient dictionary (10) from the best-response function (7) via a first-order linearization of the symmetric equilibrium around the sample mean of BankExposure.

C.1 Symmetric-cell linearization

Consider a representative bank j in cell c. Peers share the cell-level private signal up to small variation around the cell mean, and symmetric reasoning applies to each peer with its own signal. Linearizing the peer equilibrium around $\lambda_1 = 0$,

$$\bar{a}_{-j}(\theta_{\text{peer}}) \;=\; \frac{\theta_{\text{peer}} - c}{\gamma - \lambda_0} \;+\; O(\lambda_1), \tag{23}$$

where θ_{peer} denotes the representative peer signal in the cell, equal to $f(\overline{\text{BankExposure}}_{-j,c}) \cdot \text{HighRisk}_c$ and approximated to first order by $\phi\,(\text{HighRisk}_c \cdot \text{PeerExposure}_{-j,c})$ around the sample mean of BankExposure. Equation (23) closes the identification of the peer-response slope: the loading on HighRisk $\cdot$ PeerExposure is not free but equals $\phi/(\gamma - \lambda_0)$ in the linearized equilibrium.

C.2 Substitution into the bank j best response

Substituting (23) into (7) and expanding to first order in λ_1,

$$\begin{aligned} a_j^* \;&=\; \frac{\theta_j - c}{\gamma} \;+\; \frac{\lambda_0 + \lambda_1\,\theta_j}{\gamma}\,\bar{a}_{-j} \\ &=\; \frac{\theta_j - c}{\gamma} \;+\; \left(\frac{\lambda_0}{\gamma(\gamma - \lambda_0)} \;+\; \frac{\lambda_1\,\theta_j}{\gamma(\gamma - \lambda_0)}\right)(\theta_{\text{peer}} - c) \;+\; O(\lambda_1^2). \end{aligned} \tag{24}$$

Expanding and collecting terms in θ_j, θ_{peer}, and $\theta_j\,\theta_{\text{peer}}$,

$$a_j^* \;=\; \frac{1}{\gamma}\,\theta_j \;+\; \frac{\lambda_0}{\gamma(\gamma - \lambda_0)}\,\theta_{\text{peer}} \;+\; \frac{\lambda_1}{\gamma(\gamma - \lambda_0)}\,\theta_j\,\theta_{\text{peer}} \;+\; \text{const} \;+\; O(\lambda_1^2), \tag{25}$$

where "const" collects terms proportional to c that are absorbed into the bank-by-cell fixed effect in (9).

C.3 Translation to empirical regressors

Substitute $\theta_j = \phi\,(\text{HighRisk}_c \cdot \text{BankExposure}_j)$ and $\theta_{\text{peer}} = \phi\,(\text{HighRisk}_c \cdot \text{PeerExposure}_{-j,c})$ into (25). Using $\text{HighRisk}_c^2 = \text{HighRisk}_c$ for the binary indicator,

$$\begin{aligned} a_j^* \;=\; & \frac{\phi}{\gamma}\,(\text{HighRisk}_c \cdot \text{BankExposure}_j) \;+\; \frac{\phi\,\lambda_0}{\gamma(\gamma-\lambda_0)}\,(\text{HighRisk}_c \cdot \text{PeerExposure}_{-j,c}) \\ & +\; \frac{\phi^2\,\lambda_1}{\gamma(\gamma-\lambda_0)}\,(\text{HighRisk}_c \cdot \text{BankExposure}_j \cdot \text{PeerExposure}_{-j,c}) \;+\; \text{const} \;+\; O(\lambda_1^2). \end{aligned} \tag{26}$$

Because the originations regression uses $\ln(1 + \text{Limit}_{\text{new}})$ as the dependent variable and adaptation a_j corresponds to a *reduction* in originations, the empirical coefficients in (9) carry the opposite sign:

$$\beta_1 \;=\; -\frac{\phi}{\gamma}, \qquad \beta_2 \;=\; -\frac{\phi\,\lambda_0}{\gamma(\gamma-\lambda_0)}, \qquad \beta_3 \;=\; -\frac{\phi^2\,\lambda_1}{\gamma(\gamma-\lambda_0)}, \tag{27}$$

which reproduces the dictionary (10).

C.4 Why the short specification is biased

If the regression (9) is run *without* the peer triples, the coefficient on $\text{HighRisk}\cdot\text{BankExposure}_j$ picks up the partial derivative of a_j^* with respect to BankExposure_j evaluated at the sample mean of $\text{PeerExposure}_{-j,c}$. From (26),

$$\beta_1^{\text{short}} \;=\; -\frac{\phi}{\gamma} \;+\; \frac{\phi^2\,\lambda_1}{\gamma(\gamma-\lambda_0)}\,\overline{\text{PeerExposure}} \;\neq\; \beta_1. \tag{28}$$

The short-specification coefficient therefore absorbs a component of the complementarity channel; it does not identify the direct effect $-\phi/\gamma$. The full specification — in which the triple interaction is explicitly included — is the correct identification sample for structural interpretation of the own-exposure coefficient. This is the formal content of the attenuation

prediction (P4) in the main text.

C.5 Ratio invariants

Two invariants of (27) are worth noting, as they provide specification checks:

$$\frac{\beta_2}{\beta_1} = \frac{\lambda_0}{\gamma - \lambda_0}, \tag{29}$$

$$\frac{\beta_3}{\beta_2} = \frac{\phi\,\lambda_1}{\lambda_0}. \tag{30}$$

The first ratio is a pure structural object (no ϕ); the second combines the peer-norm and adverse-selection parameters with the signal slope. Stability of (29) across Y-14M and Equifax samples is a specification check: under correct specification, both samples should estimate the same λ_0/γ. The multiplier $\mathcal{M} = 1/(1 - \lambda_0/\gamma)$ that we defer in this version is obtained by inverting (29):

$$\frac{\lambda_0}{\gamma} = \frac{\beta_2/\beta_1}{1 + \beta_2/\beta_1}, \qquad \mathcal{M} = 1 + \frac{\beta_2}{\beta_1}.$$

Clean implementation of $\mathcal{M}$ in the Equifax sample requires the coefficient on HighRisk $\times$ PeerExposure in the full specification, which is not reported in the current Table 13; the full-sample Equifax quantification is left to future work.