When Banks Wait: Strategic Delay for Distressed Loan Resolution

Judy Yue*

Current draft: November 13, 2026

Update frequently, see here for the most updated version

Abstract

Why do delayed loan resolution and low trading activity in collateral markets occur together after a sectoral shock? This paper proposes a price-discovery mechanism to explain these patterns. I develop a dynamic model of loan resolution in which resolving a loan reveals information that other banks can use to obtain higher expected prices from risk-averse developers. Price discovery therefore becomes a public good. Banks have incentives to wait for others to resolve their loans, while earlier resolutions encourage more resolutions in subsequent periods. When initial uncertainty is high, these incentives can prolong distress and sustain low trading activity. I test three predictions of the model using loan-level data from the commercial real estate market following the rise of work from home. First, earlier loan resolutions in a local market predict fewer loan extensions in subsequent periods, regardless of whether the price news is good or bad. Second, for resolutions involving property sales, this relationship is significantly stronger in states where transaction prices are publicly disclosed. Third, an additional prior resolution is associated with a larger proportional increase in subsequent resolution probability in markets with greater exposure to work from home. Together, these findings support a mechanism whereby individually optimal waiting for information can lead to a persistently cold market and delayed resolution of distressed loans.

*Booth School of Business and the Kenneth C. Griffin Department of Economics, University of Chicago. Email: hyue0@chicagobooth.edu. I am indebted to my advisors, Stefan Nagel, Lars Peter Hansen, Mikhail Golosov, and Raghuram Rajan, for their invaluable guidance and support. I thank my discussant, Yunzhi Hu, for his detailed comments and suggestions. For helpful comments, suggestions, and discussions, I thank Douglas W. Diamond, Anil K. Kashyap, Fernando Alvarez, Pascal Noel, Yueran Ma, John C. Heaton, Amir Sufi, Anthony Lee Zhang, Luigi Zingales, Zhiguo He, Shohini Kundu, Maryam Farboodi, Anmol Bhandari, and Joseph L. Pagliari Jr. I also thank participants in the Asset Pricing Working Group, Economic Dynamics Working Group, and Corporate Finance Working Group at the University of Chicago, as well as participants at the Yiran Fan Memorial Conference and the Macro Finance Research Program (MFR) Summer Session for Young Scholars. This research was funded in part by the John and Serena Liew Fellowship Fund at the Fama-Miller Center for Research in Finance, University of Chicago Booth School of Business.

1 Introduction

Large declines in asset values can follow technological changes that displace existing assets or shifts in demand away from a sector. After such sectoral shocks, creditors face the choice of whether to extend or resolve distressed loans collateralized by those assets. Extension can give borrowers and creditors time to learn about the market before making further decisions. Yet it keeps credit committed to investments with large downside risks. These decisions are consequential for credit allocation, intermediaries' risk exposure, and financial stability throughout the economy.

The post-pandemic U.S. commercial real estate (CRE) market provides a vivid example of such episodes. Many industries shifted toward remote work after the advent of Zoom. This substantially weakened demand for offices and nearby retail space in central business districts (CBDs) (Gupta, Mittal, & Van Nieuwerburgh, 2026). Office transaction volume in major metropolitan statistical areas (MSAs) fell by about 64 percent between 2021 and 2023, and office price indices dropped by 43.5%. But what concerns the market the most is the impact on banks (Crosignani & Prazad, 2024). At the start of 2022, total bank-held CRE mortgage debt was $2.7 trillion, of which 71.09% was held by regional or community banks (Jiang, Matvos, Piskorski, & Seru, 2025). Banks reacted to this large price decline by continuing to extend loan maturities. Extensions remain common even among large, generally well-capitalized banks (Glancy, 2026). This resulted in a moving maturity wall for the CRE loans amid a cold transaction market for assets.

Similar patterns of limited trading and delayed loan resolution showed up in other historical episodes. In the mid-2010s, we saw a drastic decrease in taxi-medallion values following the expansion of Uber and Lyft. The taxi-medallion market became illiquid, while lenders restructured loans by reducing interest rates and extending repayment periods. Another example is the shipping downturn after the 2008 global financial crisis. We observed reduced trading in secondhand vessels alongside continued loan extensions and restructuring by banks holding shipping mortgages.

Why do delayed loan resolutions and low transaction activity occur together following large sectoral shocks? What trade-offs do banks face when deciding whether to extend or resolve distressed loans? This paper provides an explanation based on price discovery through social learning. After a one-time aggregate shock, lenders and potential buyers face uncertainty about the value of the affected collateral. Resolving a loan generates valuation information through buyers' careful due diligence and asset workout results. This information provides a reference for other market participants and lowers the uncertainty discount. Banks, therefore, have an incentive to wait to benefit from the expected price rise from peer resolutions. When banks wait, fewer resolutions and transactions take place. Information production slows, and uncertainty persists. Thus, we observe a long period of repeated extensions, even when banks have sufficient capital and there is no information asymmetry between lenders and asset buyers.

I begin by developing a dynamic model of banks' decisions about when to resolve distressed

loans. The model starts with a sectoral shock that reduces the value of banks' loans and increases valuation uncertainty. Loan values share a common component that remains fixed after the shock but is unknown to market participants, who can learn about it from realized resolutions. In each period, a bank chooses whether to retain its loan or resolve it by selling the collateral to a competitive pool of developers who specialize in redeploying the collateral. Developers are risk averse and require compensation for valuation uncertainty through a lower purchase price. If a bank decides to wait and retain its loan, it incurs a holding cost. The cost is drawn independently across banks and periods and reflects the opportunity cost of capital tied up in distressed loans. After resolution, the collateral's realized value becomes public and provides a noisy signal about the common component of loan values. The reduction in uncertainty from resolutions leads to lower uncertainty discounts in future periods and benefits banks that have yet to resolve. Each bank, therefore, weighs the cost of retaining its loan against the prospect of selling at a smaller uncertainty discount after learning from others' resolutions. In equilibrium, banks follow a symmetric cutoff strategy for resolution. Banks with sufficiently high current holding costs resolve, while those with lower costs extend and wait.

The model generates persistent delays in loan resolution through within-period strategic substitution and across-period strategic complementarity. In each period, a bank is more likely to wait if other banks are more likely to resolve, since their resolutions increase its expected payoff from resolving in the next period. The public good nature of information, therefore, weakens resolution incentives within each period. The resulting decline in resolutions makes banks more likely to wait in the next period. This complementarity comes from two channels. First, fewer resolutions leave uncertainty high and maintain incentives to wait and learn. Second, fewer resolutions leave more banks with unresolved loans and more opportunities to learn from their future resolutions. Together, these forces allow low transaction activity and continued loan extensions to persist. The resolution process is therefore prolonged when the economy starts with greater initial uncertainty.

The model yields three predictions that guide the empirical analysis. First, earlier resolutions reduce subsequent loan extensions conditional on loan survival, controlling for market size and initial uncertainty. The result follows directly from across-period complementarity. Second, the response of loan extensions to earlier resolutions is stronger when each resolution reveals more precise information about collateral values. This is driven by the precision channel of complementarity. Third, greater initial uncertainty amplifies the proportional response of resolution probability to earlier resolutions. This prediction shows how complementarity varies with initial uncertainty, reflecting the convexity of the information gain. These predictions connect the extension decision, conditional on a loan remaining unresolved, to the number of earlier resolutions, along with measures of the informativeness of resolution outcomes and initial uncertainty.

To test these predictions, I construct a loan-level panel following bank-held CRE loans outstanding at the beginning of 2020 through 2025Q2. I combine CoreLogic mortgage filings with property

transactions, tax assessments, and bank regulatory data, and clean the loan-level histories, including key events such as modifications, refinancing, and mortgage releases. From these histories, I identify loan extensions and resolution events related to public information disclosure. I define an extension as continued financing by the incumbent bank with a later maturity date. Resolution events are restricted to two types that match the model's concept of information-generating events. The first is default, defined as a transfer of ownership from the borrower to the creditor. The second is a sale-payoff event, defined as loan repayment associated with a property sale. Since the predictions concern decisions across periods, I divide the data into two fixed windows. The earlier resolution window covers 2020Q1–2021Q4, and the subsequent outcome window covers 2022Q1–2025Q2. The regression panel tracks loans outstanding at the start of 2022Q1 during the second window. It contains 626,839 loans across 2,089 counties, held by 2,900 banks. The key explanatory variable is the number of earlier resolutions, including defaults and sale-payoffs, in each county–property-type market. The outcome variable is the loan-level extension indicator in the second window.

Based on this panel, I show that loans are less likely to be extended following more resolutions in the local market, consistent with the first prediction. In the OLS specification, a one-standard-deviation increase in earlier resolutions within the same county–property-type market is associated with a 0.264 percentage-point reduction in subsequent extension probability, equivalent to roughly one-tenth of the mean extension probability in the panel. The specification includes controls for initial market size, the earlier resolution rate, loan characteristics, and market-level measures of work-from-home (WFH) shock exposure, property market liquidity, and lending concentration. It also includes county, property-type, and parent-bank fixed effects.

Yet this OLS estimate is still subject to endogeneity concerns. The correlation between earlier resolutions and subsequent extensions can arise from a persistent local shock, such as unobserved perceptions about tenants' recovery. To address this concern, I construct shift-share instruments for resolution events following Greenstone, Mas, and Nguyen (2020). Intuitively, earlier resolutions can be driven by bank-wide pressures, such as capital buffer concerns or risk exposure during COVID. Local resolution counts and rates can therefore depend on the market's exposure to different banks. I construct the shifters from bank resolution propensities using resolutions outside each market, accounting for market-level fixed effects. The shares are banks' initial local loan shares. Combining the shifters and shares, I construct predicted resolution rates and counts. The predicted measures are higher where banks that resolve more elsewhere initially held larger local shares. Identification relies on the exogeneity of banks' initial local loan shares, conditional on the controls and fixed effects. The 2SLS estimate of the extension response to a one-standard-deviation increase in earlier resolutions is -0.229, which is slightly smaller but very similar in magnitude to the OLS estimate.

Next, I examine the second prediction by showing that the negative correlation between subsequent extensions and earlier resolutions is stronger when prices associated with resolution events are publicly disclosed. Variation in resolution informativeness naturally maps to variation in

state-level property price disclosure policies. In over 10 U.S. states, property transaction prices are not required to be publicly disclosed. Unlike in disclosure states, where sale prices become publicly accessible, lenders and other investors in nondisclosure states have more limited public access to previous transaction prices for valuing similar collateral. This naturally makes sale-payoff events less informative, as a piece of key information from loan repayment after a property sale is the transaction price. I therefore restrict resolution counts to sale-payoff events and show that the estimated reduction in extension probability associated with a given increase in sale payoffs is nearly twice as large in disclosure states as in nondisclosure states. As an additional comparison, I jointly include sale payoffs, defaults, and their interactions with disclosure status. Since defaults do not necessarily reveal an arm's-length sale price, they can serve as a comparison for the sale-payoff results. Consistent with the prediction, the sale-payoff interaction remains negative and statistically significant, whereas the default–disclosure interaction is statistically insignificant and has a positive point estimate. These results support the model's prediction that resolution informativeness can affect the strength of complementarity.

Finally, I test the third prediction by examining whether greater initial valuation uncertainty strengthens lenders' responses to earlier resolutions. I proxy variation in initial uncertainty with predetermined exposure to working from home, constructed from the pre-pandemic industry composition of tenants and industry-level feasibility of remote work. Greater exposure to remote work among tenants in the local market can lead to higher uncertainty about future rents, occupancy, and transaction prices for CRE collateral, resulting in larger uncertainty discounts when lenders seek to resolve their distressed loans. I estimate an exponential conditional mean of the resolution dummy using Poisson pseudo-maximum likelihood to align with the prediction about proportional changes in resolution probability. In the regression, I interact earlier resolution counts with WFH exposure, which yields a positive and statistically significant interaction coefficient across specifications. That is, when a market is more exposed to the WFH shock, earlier resolutions are associated with a larger proportional increase in subsequent resolution probability.

The three tests provide evidence consistent with the model's predictions. One concern when interpreting these findings is an alternative mechanism: what if the correlations we see in the data are driven by regulators' tools for detecting banks' delayed loss recognition? More precisely, if we assume regulators evaluate banks' extension decisions using comparable prices, then it is harder for banks to justify extensions when there are more resolutions that reveal low comparable prices for their collateral. This can also potentially generate the correlations shown. However, under this interpretation, extensions should decline primarily after earlier resolutions that reveal low asset valuations. This is contrary to my model's prediction that extensions respond to all resolutions regardless of their realized outcomes, as both good and bad valuations offer more information to help banks and developers make decisions. To examine this distinction, I design a measure of valuation news based on the price information associated with transactions. I focus on sale-payoff events and restrict the sample to disclosure states. I then measure price news as

the difference between the log transaction price and the property's assessed-value benchmark before the transaction. I classify relatively "good" and "bad" price news according to whether the price gap is above or below the pooled mean. When I include both types of events in the regression, the results show that a one-standard-deviation increase in sale-payoff events with either "good" or "bad" price news is associated with a 0.67–0.69 percentage-point reduction in extension probability, and both coefficients are statistically significant. Additionally, the coefficient on total sale-payoff events remains largely unchanged after controlling for average price news, and we see an insignificant coefficient on price news. This finding aligns well with the model's feature that declines in extensions do not differ across signal realizations.

This paper offers an explanation for delayed loan resolution that complements the risk-shifting explanation of zombie lending. In the model, waiting for information can sustain loan extensions even when banks are well capitalized and there is no information asymmetry between lenders and buyers. The paper also provides new empirical tests of dynamic social learning. Existing research documents information spillovers across firms and incentives to delay investment while waiting to learn from others. I extend this work to the resolution of distressed credit, where individual decisions affect the public information available for valuing other outstanding loans. The empirical contribution is to combine tests based on public access to transaction prices, initial uncertainty, and realized price news. These tests help distinguish the information channel from other explanations for spillovers across lenders. Together, the model and evidence connect the public-good nature of price discovery to the persistence of financial distress after sectoral shocks.

This mechanism has policy implications for the post-pandemic CRE downturn and beyond. For loans impaired by a sectoral shock, uncertainty about collateral values can further depress sale prices and delay resolution. The appropriate policy response depends on the source of the delayed resolution. If weakly capitalized banks extend loans to postpone loss recognition or shift risk, supervisors may focus on these banks and provide capital support to facilitate resolution. If the delay is driven by fire-sale pressure from funding constraints, providing capital to buyers may facilitate the transactions. The delay from social learning creates an additional consideration. Each bank may prefer to wait and learn from others' resolutions without fully internalizing the informational benefits that its own resolution would provide to other lenders. Subsidizing early resolutions could, therefore, help generate public information and encourage subsequent resolutions. Intervention timing matters because information generated early can inform resolution decisions while many loans remain outstanding. A system or platform that collects and publicly disseminates information on distressed property sales could complement this support by making the information generated by early transactions readily available to other lenders.

Related Literature This paper relates to three strands of literature. First, I contribute to the literature on zombie lending and loan forbearance. Caballero, Hoshi, and Kashyap (2008) document

how lending to insolvent firms impeded restructuring and depressed investment and employment at healthy firms in Japan. Bruche and Llobet (2014) explain this zombie-lending behavior through insolvent banks' incentives to conceal losses and gamble for recovery under limited liability. Hu and Varas (2021) show how banks' private information about borrowers can also induce them to roll over bad loans in anticipation of repayment through future market financing. I identify an additional incentive for continuation that arises when banks and potential collateral buyers share the same information about collateral fundamentals. The mechanism connects continued lending to banks' incentives to wait for public information and makes their resolution decisions interdependent.

The paper also contributes to the heated discussion on post-pandemic commercial real estate and bank lending. Gupta et al. (2026) document the effects of remote work on office income and valuations. Using supervisory data on large banks, Crosignani and Prazad (2024) link weaker capitalization to extensions of distressed CRE loans. Using the same supervisory data source, Glancy (2026) documents that extensions are also common among generally well-capitalized banks and examines their terms and subsequent performance. County mortgage records provide broader coverage of banks' CRE risk exposure beyond the largest banks. Anenberg, Kim, and Moszkowski (2025) study how banks' exposure to remote work affects credit supply, while Hinzen, Severino, and Van Nieuwerburgh (2026) measure collateral distress and examine regional banks' refinancing decisions and borrower equity requirements. I extend this empirical work by reconstructing the timing and form of loan continuation and resolution. The linked loan histories allow me to test the model's predictions by examining whether earlier resolutions at other banks predict subsequent extension decisions and whether this relationship varies with the public availability and content of transaction prices. The results provide evidence on how loan resolution unfolds across large, regional, and small community banks.

Finally, the model builds on the long literature on social learning and delayed adjustment, which is widely discussed in investment theory. Chamley and Gale (1994) show how investors delay decisions to learn from private information revealed by others' actions. Veldkamp (2005) explains slow recoveries and sudden downturns through an information flow that depends on economic activity. Fajgelbaum, Schaal, and Taschereau-Dumouchel (2017) show how feedback between uncertainty, investment, and information production can sustain low activity. I develop this logic in a model of distressed-loan resolution, where public learning reduces the uncertainty discount at which collateral can be sold. A completed resolution improves public information while removing a potential source of future information. I then show the relevance of this mechanism with a novel set of empirical designs in the post-pandemic CRE market.

The remainder of the paper proceeds as follows. Section 2 describes CRE lending and the market patterns that motivate the analysis. Section 3 develops the model and derives its empirical predictions. Section 4 describes the data and variable construction. Section 5 presents the empirical

design and results. Section 6 concludes.

2 Institutional Background and Motivating Evidence

This section provides institutional background on CRE lending. I present stylized facts about post-WFH CRE market dynamics, including the large decline in trading volume and an increase in the share of extensions. I then turn to investors' views on valuation uncertainty conveyed in industry reports. These observations motivate the model in Section 3, in which lenders delay resolution while waiting for information from other lenders' transactions.

2.1 CRE Lending and Loan Resolution

CRE is one of the major asset classes held by banks, and banks are key participants in the CRE credit market. Crosignani and Prazad (2024) report that banks held approximately half of outstanding U.S. CRE mortgage debt in 2023Q4. Throughout the paper, I focus on bank-held loans, as they share similar traits and banks' decisions are easier to track than those of other creditors.

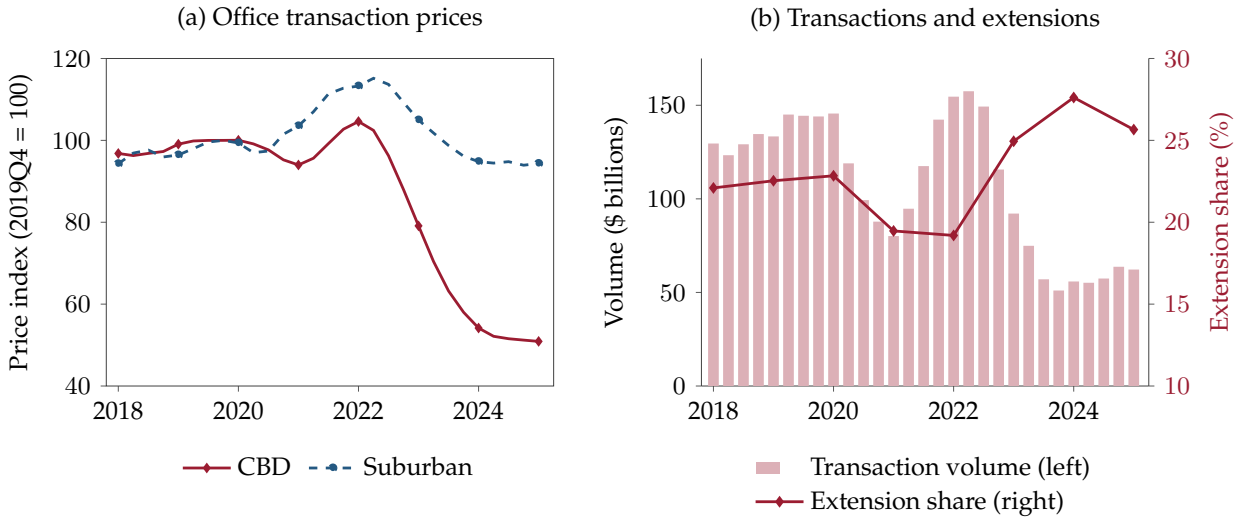
Distress in CRE loans can arise from either missed interest payments or difficulty making the balloon payment. CRE loans often mature before full amortization, leaving substantial principal due at maturity. A loan can therefore remain current on its payments yet face repayment difficulties at maturity because of decreasing property values or tighter lending conditions.

When repayment becomes difficult, the lender needs to decide whether to continue the loan or pursue resolution. An extension, through modification or refinancing with the same creditor, postpones repayment.[1] Alternatively, the borrower may repay the loan through external refinancing, a property sale, or an equity contribution. If no agreement is reached, the lender may pursue foreclosure and take control of the collateral.

An extension can be a rational choice for a bank seeking a higher payoff, even when the bank and its supervisors recognize the risks of the loan. Regulatory guidance supports prudent workouts while requiring continued risk assessment and loss recognition (Federal Reserve Board, 2023); banks also report certain troubled modifications in Call Reports (Federal Financial Institutions Examination Council, 2025). Extension gives borrowers time to restore occupancy, arrange financing, or contribute equity. An extension may also allow the lender to learn more about the property's value under a chaotic market before deciding whether to continue or resolve the loan. However, postponing principal repayment keeps the loan on the bank's balance sheet, prolonging the bank's exposure to credit risk and tying up funds that could be invested in other profitable sectors.

[1] For institutional details on the types of CRE loan modifications, see Glancy, Kurtzman, and Loewenstein (2022), Section 2.1 and Appendix A.

Figure 1: Office Prices, Transactions, and Loan Extensions



Notes. Panel A shows national office price indices for central business districts (CBD) and suburban locations, normalized to 2019Q4 = 100. Panel B shows office transaction volume as 12-month totals observed quarterly, in billions of dollars, and the annual extension share: extensions divided by extensions, refinancings, and defaults. Sources: MSCI/RCA, CoreLogic, and FFIEC.

2.2 Post-Pandemic Transactions and Loan Extensions

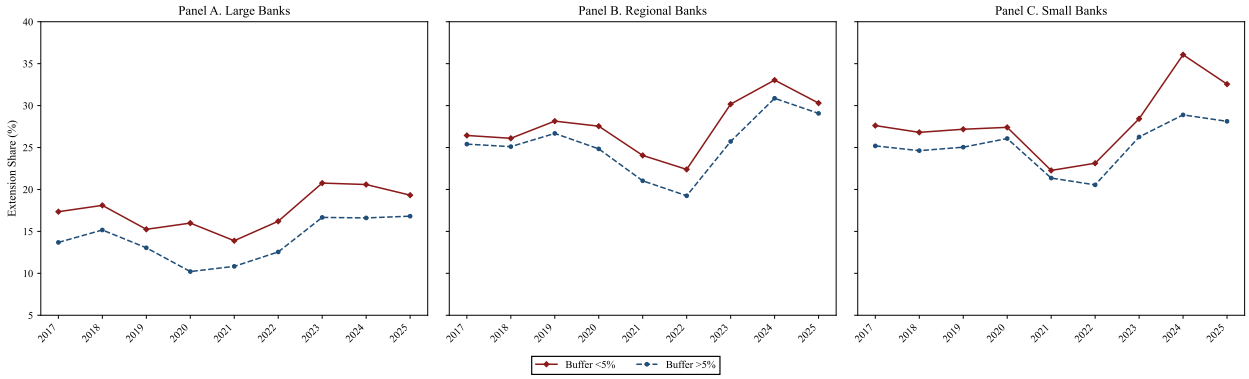
The post-pandemic office market provides the exact setting for market reaction after a lasting change in property demand. Remote work stabilized well above its pre-pandemic level (Barrero, Bloom, & Davis, 2023), while office leasing income declined. In the CompStak data, total office leasing revenue in 2025Q1 remained well below its December 2019 level across geographic markets (Appendix Figure 14). Using a valuation model incorporating cash flows and discount rates, Gupta et al. (2026) estimate that long-run New York City office values are approximately 46 percent below their pre-pandemic level. As loans approach maturity and existing leases expire, lenders face extension and repayment decisions under conditions substantially different from those before the pandemic.

Office transaction prices were still adjusting several years after the pandemic. Panel A of Figure 1 shows the MSCI/RCA office price indices rebounding in 2021 and declining again from 2022. By 2025Q1, central business district and suburban office indices were respectively 49.1 and 5.6 percent below their 2019Q4 levels.

The contraction in office transactions after 2022 coincided with greater reliance on loan extensions. Panel B of Figure 1 shows office transaction volume falling to $51.0 billion in 2023 from $144.0 billion in 2019. I define the extension share as the number of extensions divided by the total number of recorded extensions, refinancings, and defaults. This share increased from approximately 19 percent in 2022 to 28 percent in 2024.

The increase in extension shares was widespread across lenders, including banks with strong pre-

Figure 2: Office Loan Extensions by Bank Size and Capital Buffer



Notes. The extension share is the number of extensions divided by extensions, refinancings, and defaults. The sample comprises recorded outcomes for bank-held office loans, excluding CMBS-linked collateral. The large, regional, and small banks are defined based on asset cutoffs of $100 billion and $1.284 billion. Capital buffer groups use the 2019Q4 mark-to-market CET1 buffer. Sources: CoreLogic and FFIEC.

pandemic capital buffers. Figure 2 groups banks by size and their 2019Q4 mark-to-market common equity Tier 1 (CET1) buffers. Extension shares rose after 2022 on both sides of the cutoff in each size group. Glancy (2026) documents extensions at banks comfortably above capital requirements and small, statistically insignificant differences in extension-rate changes across buffer groups. Hinzen et al. (2026) find that regional banks' leniency in equity contributions at refinancing does not vary with capitalization.

Conventional explanations attribute delayed loan resolution to banks' incentives to postpone loss recognition and gamble on valuation increases because of risk-shifting (Bruche & Llobet, 2014; Caballero et al., 2008). These incentives are mostly relevant for weakly capitalized lenders, as they face a positive probability of default. For well-capitalized banks, these incentives do not readily explain extension behavior. Thus, the broader increase in extensions motivates an additional channel.

This crisis follows a change in asset fundamentals caused by an unprecedented demand shift. To explain both widespread extension increases and limited transactions, I turn to uncertainty and learning, drawing on the industry reports and comments below.

2.3 Valuation Uncertainty and Price Discovery

Apart from the trends and numbers, how do market participants understand the office market? Industry reports highlight the difficulty of valuing properties when comparable transactions are scarce. During its fourth-quarter 2023 earnings call, BXP, a large office real estate investment trust (REIT), discussed opportunities to buy properties from owners and lenders facing financial pressure. Asked what returns its investment partners would require, CEO Owen Thomas said

> "In terms of pricing, I think it's to be determined. I think it depends a lot on the building, leasing status and what exactly the issues are. There's no doubt that pricing has changed."

A follow-up question asked how office capitalization rates were changing and when they would stabilize. Thomas responded

> "You're asking a very good but unanswerable question, and I'll explain why. We tried. I try every quarter to give all of you comparable market deals. So we all understand together what pricing is. And this quarter, I could not come up with one."

The exchange illustrates how difficult it is for experienced office investors to value properties when comparable transactions are scarce. CBRE's H1 2024 Cap Rate Survey of more than 250 professionals linked limited investment volume to pricing uncertainty and wider ranges of estimated office capitalization rates (CBRE, 2024). S&P Global's January 2024 report likewise described limited price discovery alongside transaction-volume declines exceeding 60 percent in most major markets (Hudgins, 2024).

These common opinions in the industry reports suggest that investors recognized a substantial decline in office values but remained uncertain about the values of individual properties. Investors could therefore be interested in distressed properties while remaining concerned about the risks.

2.4 Link Valuation Difficulty with Loan Resolutions

Taken together, these stylized facts motivate a model in which a lasting sectoral shock lowers collateral values and leaves lenders and buyers uncertain about values under the new demand conditions. The uncertainty discount in buyers' valuations reflects their difficulty in pricing distressed properties. Completed resolutions generate public information, consistent with investors' reliance on comparable transactions. This leads to a free-riding incentive for banks to wait and benefit from the decrease in the uncertainty discount from others' resolutions. When many banks wait, fewer transactions produce new information, and uncertainty persists. This feedback provides a possible explanation for the coexistence of limited trading and widespread loan extensions we see in this section.

This mechanism complements risk-shifting explanations by showing how information free-riding can delay resolution even at well-capitalized banks. This channel can also apply beyond CRE. Following the expansion of Uber and Lyft, declining taxi income and medallion values impaired refinancing, and lenders reduced interest rates and extended repayment periods (BankUnited, Inc., 2016; Peters, 2017). In January 2018, Aspire Credit Union used auction bids to assess buyer demand and announced another auction with financing to test the market (Birch, 2018). This

episode illustrates how transactions can provide information relevant to lenders' decisions about whether to continue or resolve distressed loans.

The next section presents the model, explains its key assumptions, and examines the forces that delay resolution. I then test the model's predictions using loan-level data from the post-pandemic CRE market, examining whether earlier resolutions at other banks in the same local market are followed by fewer loan extensions. The CRE market provides rich cross-sectional variation in access to information and detailed data on collateral values and transaction prices, making it a useful setting for assessing whether this relationship reflects learning.

3 Model

3.1 Environment, Timing, and Beliefs

I develop a discrete-time model of loan resolution in which banks can wait to learn from other banks' resolutions. Time is indexed by $t = 0, 1, 2, \ldots$. There are two types of agents: banks, which hold loans backed by collateral, and developers, which specialize in realizing collateral value.

Initially, $N_0 = \bar{N} \in \mathbb{N}$ banks each hold one loan backed by one unit of collateral in a given sector, such as office buildings or taxi medallions. At $t = 0$, an aggregate demand shock sharply reduces collateral values and the value of the loans backed by them. This value has two components: θ, which is common to all the loans in the sector, and ε_i, reflecting the loan terms and characteristics of the underlying collateral. The idiosyncratic components are independently and identically distributed across assets:

$$\text{Asset Value}_i = \theta + \varepsilon_i, \quad \varepsilon_i \overset{i.i.d.}{\backsim} N(0, \delta^{-1})$$

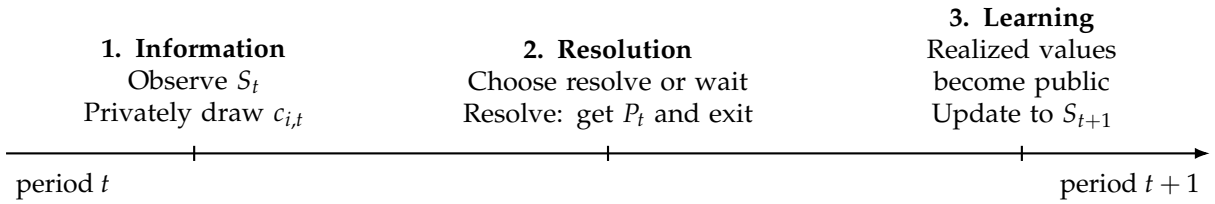
Neither banks nor developers observe θ or ε_i before resolution. The object of learning is the common component θ, which can be interpreted as the sector's post-shock fundamental. For simplicity, I assume it remains fixed after the shock. All market participants begin with the same information set $\mathcal{I}_0$ and share the same initial belief, with mean μ_0 and precision $\gamma_0 > 0$:

$$\theta | \mathcal{I}_0 \sim N(\mu_0, \gamma_0^{-1})$$

After the shock, each bank decides in every period t whether to resolve its loan or wait. Resolution ends the game for the bank and delivers the payoff P_t from developer, which can be interpreted as loan sales, repayment, or completed default workouts. A bank that waits pays a privately observed holding cost $c_{i,t}$, drawn independently across banks and periods:

$$c_{i,t} \overset{iid.}{\backsim} G$$

Figure 3: Three stages within period t



Notes. The resolution price $P_t = P(\mu_t, \gamma_t)$ uses beginning-of-period information; payment precedes the release of current resolution signals.

Here G is continuous with strictly positive density on $[\underline{c}, \bar{c}]$, where $0 < \underline{c} < \bar{c} < \infty$. Costs drawn are independent of θ and all asset-specific shocks. The idiosyncratic component ε_i is also independent of θ, with meaningful variances $\delta > 0$.

Resolution is the only source of new public information: it reveals information about collateral values, summarized by the public signal

$$x_i = \theta + \varepsilon_i$$

The idiosyncratic component has variance δ^{-1}, so each signal has precision δ. When n banks resolve, their signals provide combined precision $n\delta$.

Let $\mathcal{N}_t$ denote the set of banks whose loans remain unresolved at the beginning of period t, and let $N_t = |\mathcal{N}_t|$ denote their number. Write $d_{i,t} = 1$ for resolution and $d_{i,t} = 0$ for waiting, and let $n_t = \sum_{i \in \mathcal{N}_t} d_{i,t}$ count total resolutions in period t. The public state is $S_t = \{\mu_t, \gamma_t, N_t\}$: the posterior mean and precision, and the number of unresolved loans. Each period has three stages: (i) banks and developers observe S_t, and each unresolved bank privately draws $c_{i,t}$; (ii) banks simultaneously choose whether to resolve or wait, and resolving banks exit; (iii) collateral-value information becomes public, beliefs update, and the economy enters period $t+1$. Figure 3 summarizes this sequence.

Realized asset values provide information about the common fundamental θ. Section 3.4 discusses this learning assumption after the pricing and resolution problems are introduced.

3.2 Developers' Demand and Resolution Prices

Uncertainty affects resolution prices through developers' risk aversion. There are $M > \bar{N}$ short-lived, competitive developers with mean-variance preferences. At a price P_t per distressed loan-recovery claim, each developer chooses how many claims to buy:

$$\max_k \quad \mathbb{E}_t \left[\sum_{i=1}^{k} (\theta + \varepsilon_i) - kP_t \right] - \frac{a}{2} Var_t \left[\sum_{i=1}^{k} (\theta + \varepsilon_i) \right]$$

Purchases are indivisible, $k \in \{0, 1, 2, \ldots\}$, and $a > 0$ measures risk aversion. Substituting the conditional means and variances into the objective, competition implies

$$\max_k k \left(\mu_t - P_t - \frac{a}{2}\frac{1}{\delta} \right) - \frac{a}{2} k^2 \frac{1}{\gamma_t} = 0$$

that is, developers earn zero risk-adjusted surplus in equilibrium. Competition among $M > \bar{N}$ developers sets the one-unit surplus to zero, giving the price below. At this price, the objective equals $ak(1-k)/(2\gamma_t)$: buying zero or one unit gives zero surplus, while buying two or more gives negative surplus.

$$P_t = \mu_t - \frac{a}{2\gamma_t} - \frac{a}{2\delta}$$

Lower public precision γ_t produces a larger uncertainty discount. The assumption $M > \bar{N}$ ensures that the developer pool can absorb all claims without a fire-sale discount. This isolates price discovery as the source of delayed resolution.

3.3 Banks' Resolution Decisions

Banks are risk neutral and do not discount future payoffs. At public state (μ_t, γ_t, N_t), a bank with waiting cost $c_{i,t}$ chooses between resolution and continuation:

$$V(\mu_t, \gamma_t, N_t, c_{i,t}) = \max \left\{ V^R(\mu_t, \gamma_t),\ V^W(\mu_t, \gamma_t, N_t, c_{i,t}) \right\}, \tag{3.1}$$

The resolution value $V^R(\mu_t, \gamma_t)$ is the current price and depends only on public beliefs. The waiting value deducts the current holding cost from the expected continuation value, conditional on the bank remaining unresolved:

$$V^R(\mu_t, \gamma_t) = P(\mu_t, \gamma_t) = \mu_t - \frac{a}{2\gamma_t} - \frac{a}{2\delta}, \tag{3.2}$$

$$V^W(\mu_t, \gamma_t, N_t, c_{i,t}) = -c_{i,t} + \mathbb{E}_t \left[V(\mu_{t+1}, \gamma_{t+1}, N_{t+1}, c_{i,t+1}) \right]. \tag{3.3}$$

For compactness, write $V(S, c)$, $V^R(S)$, and $V^W(S, c)$ for these values. Before its cost draw, the bank's ex ante value is

$$\mathcal{V}(S) = \int V(S, c)\, dG(c).$$

This value depends only on the public state. The following lemma shows that the posterior mean shifts the bank's value without changing its resolution decision.

Lemma 1. The ex-ante value takes the form

$$\mathcal{V}(\mu_t, \gamma_t, N_t) \equiv \int V(\mu_t, \gamma_t, N_t, c_{i,t})\, dG(c_{i,t}) = \mu_t + v(\gamma_t, N_t).$$

Consequently, the equilibrium cutoff and resolution probability do not depend on μ.

Proof: See Appendix A.4 for the full proof. The posterior mean is a martingale, and both resolution and continuation payoffs are affine in that mean. Subtracting it from the Bellman equation leaves a problem in (γ, N), so the cutoff and resolution probability are independent of μ.

Learning from others' actions is central to Chamley and Gale (1994), while Kurlat (2018) studies how a trading history helps agents value assets. Here, realized signals change the posterior mean, but that mean shifts the risk-neutral bank's resolution and continuation values equally. Its resolution decision therefore depends on public precision and the number of unresolved banks, rather than on whether the realized news raises or lowers μ.

3.4 Discussion of the Assumptions

The model's learning mechanism rests on three ingredients: asset values share an uncertain component, completed resolutions reveal public information about that component, and uncertainty lowers the payoff from resolution. I discuss their economic roles and relevance to CRE. I then turn to the assumptions introduced to simplify the equilibrium analysis.

3.4.1 Common Uncertainty and Public Learning

A common source of uncertainty makes information from one resolution useful to other banks. The model captures this through θ: each asset's value depends on the same unobserved fundamental, so observing others' realized values helps banks assess the value of their own assets. This source of learning also remains with known exposure $\alpha \neq 0$ in Asset Value $= \alpha\theta + \varepsilon_i$.

Linking to the case of the post pandemic CRE market, the work-from-home (WFH) shock illustrates how loan values reflect both common uncertainty and loan-specific variation. For example, firms' shift to WFH affects demand for office space as well as nearby retail, industrial, and hotel properties. Uncertainty about demand after the shock corresponds to the common component θ, while properties' exposure varies with the industry composition of local tenants. The demand uncertainty is not just restricted to the current tenant preference, but also to the residual value of the assets, e.g., which asset can be converted to other uses, and whether we can expect a return of tenants after renovations. Under this uncertainty, a conversion/renovation plan or a workout strategy, resulting in either a good or bad signal for asset valuation, can reveal the recovery values under these demand conditions and lower the risks for the developers. The idiosyncratic noises can come from property characteristics and existing leases, which limit what one resolution reveals about other assets.

The second ingredient is that resolution makes new information about collateral values available

to other banks. In CRE, this includes workout outcomes, conversion progress, vacancy rates after conversion, and publicly disclosed collateral sale prices. In the model, x_i denotes the public signal generated by resolution, while P_t denotes the bank's payoff from resolving.

Thus, the key parameters that determine the scope for public learning are the source of uncertainty and the informativeness of each resolution result. Informational spillovers disappear when asset values are entirely idiosyncratic. A temporary liquidity shock offers less scope for this learning than common uncertainty about fundamentals. Private workout outcomes, silent loan transfers, or refinancing without revealed collateral information contribute relatively little public information and thus provide less incentive for waiting. CRE distress after a permanent demand shift combines common uncertainty with observable valuations and workout outcomes.

3.4.2 The Uncertainty Discount

The third ingredient is that uncertainty lowers the bank's payoff from resolution. Risk-averse developers generate this discount in the model, capturing buyers' limited risk-bearing capacity or financing constraints. Holding expected asset value fixed, greater uncertainty lowers their offers.

Through this discount, uncertainty enters banks' decisions as a key state variable in general equilibrium. Banks' resolution outcomes reveal information about the common fundamental that can reduce the discount demanded by future buyers. Since banks are risk neutral, there are zero utility costs of the risks of how the same information changes the μ. This simplifies the tradeoff for banks, leaving them to weigh the information benefit against the holding costs. With risk-averse banks, bearing this uncertainty creates an additional cost in the banks' problem. A sufficient condition for generating waiting is that the buyers are relatively more risk-averse than banks.

The assumption of risk-averse buyers has an institutional rationale. Buyers of distressed assets, such as specialized developers financed by closed-end funds, can face higher costs of bearing property risk than well-capitalized banks holding the existing loans. Developers' concentrated property exposure limits diversification (Ellis & Naughtin, 2010). Closed-end funds have finite horizons, and interim performance affects subsequent fundraising (Barber & Yasuda, 2017). On the contrary, banks' deposit franchises can support stable funding and longer-term lending (Drechsler, Savov, & Schnabl, 2021), while buyers seeking new financing face lending conditions that tighten with short-term real estate uncertainty (Cao, Song, Titman, & Zhan, 2025) .

I can also relax the assumption of risk-averse buyers. This changes the interpretation but preserves the logic through which learning generates waiting. The conceptual requirement for the uncertainty-discount channel is that, holding the expected asset value fixed, greater common uncertainty lowers the bank's risk-adjusted payoff from resolution. This can also arise in an extension with risk-averse banks and risk-neutral property buyers. Consider a bank choosing between foreclosure with uncertain proceeds from a future auction and a settlement offering a

known recovery. Suppose the bank cannot sell its foreclosure claim upfront, and both options remain available while it waits. The bank then bears the risk in foreclosure proceeds. A mean-preserving spread in these proceeds due to greater common uncertainty lowers their certainty equivalent. Other banks' sales can reveal common market conditions before the bank commits to a resolution method. The bank can use this information to choose foreclosure when its risk-adjusted payoff exceeds the settlement recovery and accept the settlement otherwise. If the benefit of making a better resolution decision outweighs the costs and risks of waiting, the bank has an incentive to postpone resolution to learn from others. This extension therefore provides another way for learning to generate delay. For tractability and to reflect the institutional setting, the baseline models the waiting incentive through buyers' uncertainty discounts.

3.4.3 Simplifying Assumptions

The remaining assumptions isolate price discovery and keep the resolution problem tractable.

Information and payment timing. In the baseline, the bank knows the resolution payoff P_t when it decides whether to resolve. Payment could instead follow an appraisal or bargaining after irreversible commitment. Risk neutrality preserves the decision if expected payment equals P_t and the public signal process remains unchanged.

Waiting-cost heterogeneity. The positive per-period holding cost $c_{i,t}$ captures the cost of retaining a distressed loan while waiting for information. It can represent the opportunity cost of capital that could be deployed elsewhere., which can be idiosyncratic across bank and time. This setup for the private holding-cost shocks follows the logic of Harsanyi's equilibrium-purification argument (Harsanyi, 1973), enabling each agent to make a deterministic decision conditional on its own idiosyncratic draw, while uncertainty about that draw makes its choice probabilistic for other agents. Yet this assumption is only for tractability and is not necessary for generating delay. Appendix A.2 shows that the incentive to wait also arises with fixed costs.

Waiting costs are assumed to be uninformative about the common fundamental. Banks privately observe their own waiting costs but have no other private information. Thus, conditional on public information, a bank's decision to resolve or wait does not itself reveal information about θ. If banks instead held private information about collateral values, their decisions could provide an additional source of learning. These assumptions abstract from that channel and isolate learning from the outcomes revealed by resolution.

A deep pool of developers. We assume a large, competitive developer pool to remove the fire-sale force. With $M > \bar{N}$, additional resolutions do not lower the current price with fixed public beliefs.

This removes strategic interaction through current price pressure and isolates the role of learning. The assumption describes a market in which investors have capital, but wide bid–ask spreads impede transactions.

Fixed fundamentals. The fundamental remains fixed after the initial shock, while banks and developers update their beliefs as resolutions reveal information about its value. This separates waiting to learn about the post-shock value from waiting for the fundamental itself to change.

One asset per bank. I assume that each bank holds only one asset for simplicity, but the incentive to delay resolution can persist when a bank holds several assets. Resolving one loan can inform decisions about the rest of its portfolio. A bank may still wait when learning from other banks outweighs holding costs and the benefit of generating its own information. The one-asset assumption removes this additional benefit of learning within the bank's portfolio, allowing the baseline to isolate delays driven by learning from others.

Together, these assumptions isolate the tradeoff between banks' costs of waiting and the potential information gains from other banks' resolutions.

3.5 Equilibrium

I now characterize equilibrium in the resolution game. I first define a symmetric Markov equilibrium, then derive its cutoff representation and establish existence and uniqueness.

3.5.1 Equilibrium Definition

An unresolved bank chooses whether to resolve using the public state $S_t = (\mu_t, \gamma_t, N_t)$ and its current waiting cost c. Costs are independent across banks and periods and independent of θ, so (S, c) is the bank's payoff-relevant state. I consider symmetric Markov equilibria, in which all unresolved banks use the same measurable strategy

$$s : \mathbb{S} \times [\underline{c}, \bar{c}] \to \{0, 1\}, \qquad \mathbb{S} \equiv \mathbb{R} \times (0, \infty) \times \{1, \ldots, \bar{N}\},$$

Here $s(S, c) = 1$ denotes resolution. When all other unresolved banks follow s, each resolves independently with probability

$$q^s(S) = \int \mathbf{1}\{s(S, c) = 1\}\, dG(c), \tag{3.4}$$

A bank that waits at S therefore faces a binomial number of peer resolutions:

$$n_{-i} \sim \text{Binomial}(N-1,\, q^s(S))$$

where n_{-i} counts peer resolutions, conditional on bank i waiting. A bank receives $P(\mu,\gamma) = \mu - a/(2\gamma) - a/(2\delta)$ upon resolution and pays c upon extension. For $n \geq 1$ total resolutions, let $X_t(n)$ be the average of the n publicly revealed signals. The public state evolves by Bayesian updating:

$$\mu_{t+1} = \frac{\gamma_t \mu_t + n\delta X_t(n)}{\gamma_t + n\delta}, \qquad \gamma_{t+1} = \gamma_t + n\delta, \qquad N_{t+1} = N_t - n$$

If no bank resolves, the public state is unchanged. Under a symmetric strategy, the total count satisfies $n_t \sim \text{Binomial}(N_t, q^s(S_t))$; the waiting bank's problem uses the $N_t - 1$ peers instead.

Given a symmetric strategy, equilibrium requires each bank's strategy to be optimal against the others' strategies, with public beliefs following the updating rule above. The existence theorem below establishes that equilibrium values are finite and that resolution occurs almost surely.

Definition 1. A ***symmetric Markov perfect equilibrium*** is a symmetric Markov strategy s^* such that, at every state $(S, c) \in \mathbb{S} \times [\underline{c}, \bar{c}]$, the action $s^*(S, c)$ maximizes the bank's expected payoff when all other unresolved banks follow s^* at every current and future state.

Because private costs are serially independent and uninformative about θ, the state need not include beliefs about other banks' types. Cost heterogeneity purifies the equilibrium: banks use pure cutoff strategies. Banks interact through the public state and take the resolution price as given. This structure also permits an equivalent lagged-state recursive representation based on the previous period's beliefs μ_-, γ_-, the previous number of unresolved assets N_-, and the realized resolution count n. The uncertainty arises from the realization of random cost draws among a finite number of banks, generating variation in resolution history. I present that lagged-state recursive representation more formally in Appendix A.1. Throughout the main text, S denotes the current period public state (μ, γ, N) from the symmetric Markov perfect equilibrium definition.

3.5.2 Existence and Uniqueness

Lemma 1 reduces the equilibrium problem to the mean-adjusted ex ante value $v(\gamma, N)$. In equilibrium, banks follow a symmetric cutoff strategy:

$$\text{resolve if and only if } c \geq c^*(\gamma, N),$$

The corresponding resolution probability is

$$q^*(\gamma, N) = 1 - G(c^*(\gamma, N))$$

Because N is finite, the equilibrium can be constructed recursively. For $N \geq 1$, take the following expectations over $n \sim \text{Binomial}(N-1, q^*(\gamma, N))$. The reduced value and cutoff solve

$$v(\gamma, N) = \mathbb{E}_c\left[\max\left\{-\frac{a}{2\gamma} - \frac{a}{2\delta}, -c + \mathbb{E}_n\big[v(\gamma + n\delta, N-n)\big]\right\}\right],$$
$$c^*(\gamma, N) = \frac{a}{2\gamma} + \frac{a}{2\delta} + \mathbb{E}_n\big[v(\gamma + n\delta, N-n)\big].$$

The cutoff $c^*(\gamma, N)$ is the benefit of waiting before paying the current holding cost: expected continuation value minus the current resolution payoff. A bank waits when its cost is below this benefit and resolves when its cost is at least as large. Thus, banks with higher current holding costs resolve first.

When $N = 1$, there are no peer resolutions. Because waiting is strictly costly and generates no information, the bank resolves immediately:

$$v(\gamma, 1) = -\frac{a}{2\gamma} - \frac{a}{2\delta}, \qquad c^*(\gamma, 1) = 0, \qquad q^*(\gamma, 1) = 1.$$

Thus, the state $N = 0$ is absorbing for the aggregate economy. Based on this absorbing state, we can show the monotonicity of value functions and further establish the equilibrium uniqueness.

Theorem 2. Equilibrium existence, uniqueness, and continuation-value ordering.

On $\mathbb{S}$, there exists a unique symmetric Markov cutoff equilibrium. At every active state,

$$q^*(\gamma, N) > 0.$$

Consequently, starting from any state with finitely many unresolved banks, all banks resolve in finite time almost surely.

Moreover, for every $N \geq 2$,

$$v(\gamma + \delta, N-1) > v(\gamma, N). \tag{3.5}$$

Proof: See Appendix A.5 for the full proof.

The equilibrium describes the transition from the post-shock state to one in which all assets have been redeployed by developers. Throughout this transition, the functions q^*, c^*, and v depend only on public precision γ and the number of unresolved banks N.

The ordering $v(\gamma + \delta, N-1) > v(\gamma, N)$ means that a remaining bank benefits from a completed peer resolution. It receives the information without paying the holding costs needed to wait for that resolution. The proof shows that the continuation value increases even after accounting for the loss

of one future source of information.

The general equilibrium has no closed-form solution. The next section characterizes its strategic forces and resolution dynamics. I also present a simpler fixed-cost benchmark in Appendix A.2, which illustrates the same forces with more direct algebra. Below, q^*, c^*, and v denote the equilibrium resolution probability, cutoff, and reduced value.

3.6 Strategic Waiting and Resolution Dynamics

3.6.1 Resolution Times and Baseline Dynamics

Anticipated peer information encourages banks to wait, while information from completed resolutions makes further resolution more attractive. These forces can generate slow initial resolution followed by acceleration. To illustrate the dynamics, define the resolution hazard of a surviving bank as

$$h_t = \frac{\mathbb{E}[N_t q^*(\gamma_t, N_t)]}{\mathbb{E}[N_t]} = \frac{\mathbb{E}[n_t]}{\mathbb{E}[N_t]},$$

whenever $\mathbb{E}[N_t] > 0$. This is a survivor-weighted average across states.

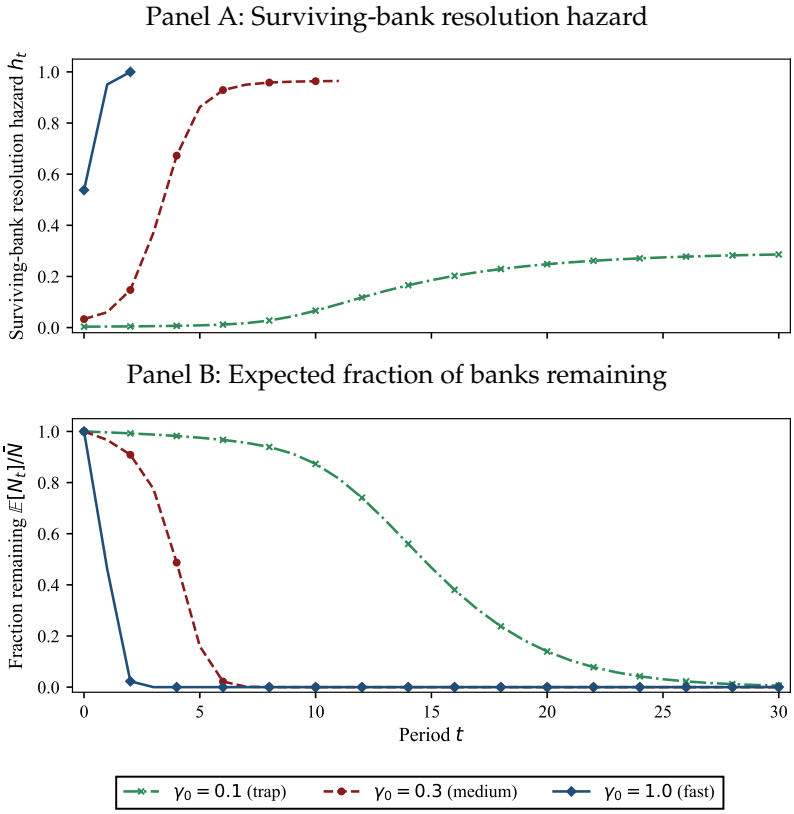
Figure 4 compares the resolution rate among surviving banks (the hazard rate) and the expected fraction of banks remaining across initial public precisions. Low initial precision produces a prolonged period of slow resolution before activity accelerates, whereas better initial information allows banks to resolve much sooner. With little initial information, banks have more to gain from waiting for peers. Few early resolutions leave public precision low; as resolutions accumulate, higher precision and fewer remaining peers encourage further resolution.

More unresolved peers and more informative future resolution signals increase banks' incentives to wait for information. Greater public precision accelerates resolution, whereas greater developer risk aversion lengthens the resolution process. Appendix A.8 states and proves these comparisons; Appendix A.3 defines and computes the expected resolution delays used in the numerical illustrations.

3.6.2 Why Is Resolution Delayed?

The delay mechanism combines within-period strategic substitution with across-period complementarity. A bank prefers to wait when its peers are more likely to resolve, because their resolutions will produce information. Once those resolutions occur, improved public information and fewer remaining peers make resolution more attractive. Without the uncertainty discount, banks would resolve immediately.

Figure 4: Numerical illustration of equilibrium resolution dynamics



Notes. Initial precision is $\gamma_0 \in \{0.1, 0.3, 1.0\}$, with $\bar{N} = 100$, $a = 2$, $\delta = 0.02$, and $G = U[0.1, 1]$. Lines are exact expectations from the equilibrium transition matrix. Panel A shows h_t; Panel B shows the expected fraction remaining. The discrete sum of the lower curve from $t = 1$ to infinity equals the expected time to resolution, with immediate resolution assigned zero delay (Appendix A.3). Hazards are omitted once the probability of an active market falls below 10^{-8}.

Theorem 3. Within-period strategic substitutes. Fix a state (γ, N) with $N \geq 2$, and consider an unresolved bank i. Hold fixed the value of continued play, and let $\widehat{q} \in [0, 1]$ denote the current resolution probability of each of the other $N - 1$ banks. Bank i's cutoff satisfies

$$\frac{\partial c_i^*(\gamma, N; \widehat{q})}{\partial \widehat{q}} > 0, \tag{3.6}$$

Thus, a higher peer-resolution probability raises bank i's resolution cutoff.

Proof: See Appendix A.6 for the full proof.

To express this incentive in terms of resolution probabilities, define bank i's best-response resolution probability,

$$q_i^{\mathrm{BR}}(\gamma, N; \widehat{q}) \equiv 1 - G\big(c_i^*(\gamma, N; \widehat{q})\big),$$

which is weakly decreasing in $\widehat{q}$ because the cutoff rises with the probability of peer resolution.

Whenever the cutoff lies in the interior of the support of G,

$$\frac{\partial q_i^{\mathrm{BR}}(\gamma, N; \widehat{q})}{\partial \widehat{q}} < 0. \tag{3.7}$$

At a symmetric equilibrium,

$$\widehat{q} = q^*(\gamma, N; \delta) \qquad \text{and} \qquad q_i^{\mathrm{BR}}(\gamma, N; \widehat{q}) = \widehat{q}$$

A bank that resolves receives its current recovery and exits, while a bank that waits can benefit from a lower uncertainty discount after other banks resolve. This creates an incentive to free ride by letting other banks resolve first. Each bank weighs the expected information benefit against its holding cost. When this benefit is large relative to holding costs, resolution can be slow in equilibrium. Since new public information arrives only through resolution, slow resolution also slows learning. As other banks resolve, public information becomes more precise and fewer banks remain as potential sources of information, making further resolution more attractive.

We next characterize how realized peer resolutions change the incentives of the remaining banks. Fix an initial state (γ_t, N_t) with $N_t \geq 2$, and consider an unresolved bank i. Conditional on bank i remaining unresolved, suppose that n of the other banks have resolved. Define the equilibrium cutoff and resolution probability along this realized resolution path by

$$c_n^* \equiv c^*\big(\gamma_t + n\delta, N_t - n; \delta\big), \tag{3.8}$$

$$q_n^{t+1}(\gamma_t, N_t; \delta) \equiv q^*\big(\gamma_t + n\delta, N_t - n; \delta\big) = 1 - G(c_n^*), \tag{3.9}$$

for $n = 0, \ldots, N_t - 1$.

Theorem 4. Across-period complementarity. For every $n = 0, \ldots, N_t - 2$, define the marginal effect of one additional realized resolution by

$$\Delta_n q_{t+1}(\gamma_t, N_t; \delta) \equiv q_{n+1}^{t+1}(\gamma_t, N_t; \delta) - q_n^{t+1}(\gamma_t, N_t; \delta). \tag{3.10}$$

Then one additional realized resolution strictly lowers bank i's subsequent cutoff and weakly raises its resolution probability:

$$c_{n+1}^* < c_n^*, \qquad \Delta_n q_{t+1}(\gamma_t, N_t; \delta) \geq 0. \tag{3.11}$$

If $q_n^{t+1} < 1$, then

$$\Delta_n q_{t+1}(\gamma_t, N_t; \delta) > 0. \tag{3.12}$$

Thus, whenever bank i has a positive probability of waiting after n resolutions, one additional resolution strictly lowers its subsequent extension probability.

Proof: See Appendix A.7 for the full proof.

The theorem compares two histories starting from the same state (γ_t, N_t): n peers resolve in one history and $n+1$ in the other. The additional resolution raises a surviving bank's resolution probability at $t+1$ whenever its probability of waiting remains positive.

An additional resolution changes both components of the reduced state. After n resolutions, write

$$N = N_t - n, \qquad \gamma = \gamma_t + n\delta.$$

The effect of one additional resolution can then be decomposed as

$$\begin{aligned}\Delta_n q_{t+1}(\gamma_t, N_t; \delta) = \; & \underbrace{q^*(\gamma, N-1; \delta) - q^*(\gamma, N; \delta)}_{\text{fewer-unresolved-banks}} \\ & + \underbrace{q^*(\gamma+\delta, N-1; \delta) - q^*(\gamma, N-1; \delta)}_{\text{public-information}}. \end{aligned} \tag{3.13}$$

The first term is weakly positive because one fewer unresolved peer reduces bank i's opportunity to wait for future peer-generated information. The second term is also weakly positive because the information generated by the additional resolution raises current public precision. The first two parts of Proposition 10 in Appendix A.8 establish these state comparisons.

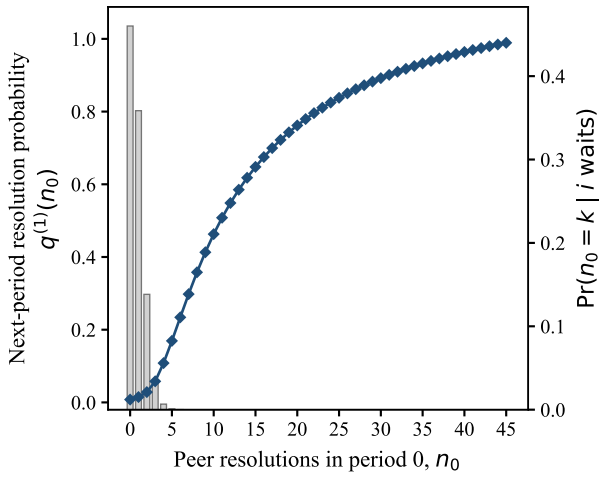
Figure 5 illustrates the response and its decomposition from a common initial state. Panel A shows the next-period resolution probability q^1 against the number of initial resolutions n_0; the background bars show the probability of each realization. Resolution probabilities rise particularly steeply between $n_0 = 0$ and $n_0 = 10$, where the initial resolution count has substantial probability mass. Panel B decomposes the effect of one additional resolution into the public-information channel (red) and the effect of one fewer unresolved peer (blue). Public information accounts for most of the increase in this calibration, which starts with many banks and low public precision. This setting is relevant to CRE distress with many lenders and limited information about affected assets' recovery values.

Together, the two strategic forces explain how resolution can remain slow before accelerating. Anticipated peer information encourages waiting. When few banks resolve, public precision remains low, and many potential sources of information remain. This gives banks an incentive to keep waiting. As banks resolve, public precision increases, and fewer banks remain as potential information sources, making subsequent resolution more attractive. Figure 4 illustrates these dynamics.

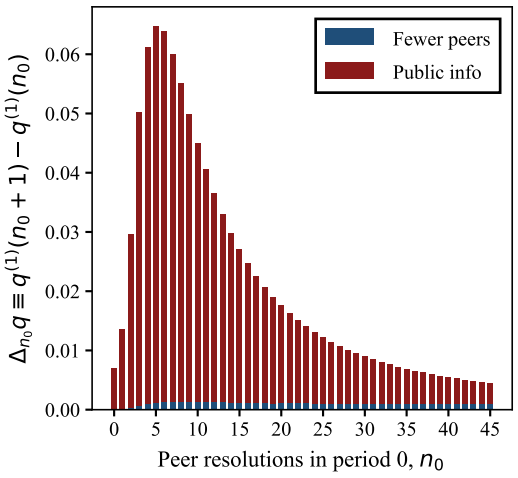
Across-period complementarity provides the main link to the empirical analysis, relating earlier resolutions to later decisions by banks that have not yet resolved. Within-period substitution is harder to examine empirically because it concerns banks' expectations about one another's simultaneous decisions, which raises a reflection problem. It is also derived by holding continuation values fixed, a condition that is difficult to capture empirically. To further explore across-period

Figure 5: Realized peer resolutions and subsequent resolution incentives

Panel A: Next-period resolution probability



Panel B: Resolution response decomposition



Notes. The line in Panel A shows $q^*(\gamma_0 + n_0\delta, \bar{N} - n_0)$; the bars (right axis) show the distribution of n_0 conditional on bank i waiting, namely Binomial$(\bar{N} - 1, q^*(\gamma_0, \bar{N}))$. Panel B decomposes the effect of one additional resolution as in (3.13). The figure's $q^{(1)}(n_0)$ denotes $q^1_{n_0}(\gamma_0, \bar{N}; \delta)$. Parameters are $\bar{N} = 100$, $\gamma_0 = 0.3$, $\delta = 0.1$, $a = 2$, and $G = U[0.1, 1]$.

complementarity, the next two results examine how the response to realized resolutions depends on the information each resolution provides and the public information already available.

3.6.3 Signal Informativeness and Amplification

Signal precision δ measures how much information a resolution supplies. I first compare the response to an additional resolution across signal precisions, holding the initial public state fixed.

The historical effect $\Delta_n q_{t+1}(\gamma_t, N_t; \delta)$ compares histories with $n + 1$ and n realized peer resolutions from the same initial state. When signal precisions differ, the public precisions after n resolutions are $\gamma_t + n\delta^H$ and $\gamma_t + n\delta^L$. The following result accounts for both this accumulated information and the information produced by the additional resolution.

Proposition 5. Signal informativeness and the amplification of across-period complementarity.

Fix $N_t \geq 3$, $\delta^H > \delta^L > 0$, and sufficiently low initial precision $\gamma_t > 0$. There exists a finite a^* such that, for every $a \geq a^*$,

$$\Delta_n q_{t+1}(\gamma_t, N_t; \delta^H) > \Delta_n q_{t+1}(\gamma_t, N_t; \delta^L), \qquad n = 0, \ldots, N_t - 3. \tag{3.14}$$

Proof: See Appendix A.9.2 for the full proof.

Appendix A.9.2 gives a sufficient condition for the proposition. Large a makes resolution sufficiently rare. The range $N_t - n \geq 3$ excludes the final comparison with a one-bank successor, which resolves with probability one. Each additional resolution raises public precision and removes a source of future learning. Under the stated conditions, the combined effect on subsequent resolution is stronger when signals are more precise.

For the first peer resolution ($n = 0$), the amplification result does not require large a; its proof appears in Appendix A.9.1. Appendix Figure 7 plots the increase in resolution probability after the first peer resolution as signal precision varies. The effect increases with δ, with a steeper slope above the sufficient-condition threshold. At very low signal precision, the curve has a small dip and is nearly flat: the information channel is weak, and changes in signal precision produce little amplification.

More precise signals strengthen the public-information channel. Intuitively, when holding the posterior mean fixed, one resolution reduces the uncertainty discount by

$$\frac{a}{2}\left(\frac{1}{\gamma_t} - \frac{1}{\gamma_t + \delta}\right)$$

And this gain clearly increases with δ. An additional force comes from future learning: a more informative resolution leaves less uncertainty for subsequent resolutions to remove. I show this relationship more directly in the Figure Appendix A.9.1, where I show the relationship in a baseline example where resolution $n = 0$. Under a sufficient precision condition, we see that more precise signals lower the pre-event resolution probability and raise the post-event probability, leading to a strictly upward-sloping curve for resolution probability response and signal precision.

3.6.4 Initial Uncertainty and Proportional Amplification

Initial public precision γ_t measures information already available, whereas δ measures the information supplied by a future resolution. I now compare economies with the same N_t and signal precision but different initial public precisions. The result concerns the proportional response to an additional resolution: a larger proportional increase can occur from a lower initial resolution probability.

Proposition 6. Initial uncertainty and proportional amplification.

Fix $N_t \geq 3$, $\delta > 0$, and $0 < \gamma^L < \gamma^H$. There exists a finite a^* such that, for every $a \geq a^*$,

$$\Delta_n \log q_{t+1}(\gamma^L, N_t; \delta) > \Delta_n \log q_{t+1}(\gamma^H, N_t; \delta), \qquad n = 0, \ldots, N_t - 3, \tag{3.15}$$

where

$$\begin{aligned} \Delta_n \log q_{t+1}(\gamma_t, N_t; \delta) \equiv\ & \log q^*(\gamma_t + (n+1)\delta, N_t - n - 1; \delta) \\ & - \log q^*(\gamma_t + n\delta, N_t - n; \delta). \end{aligned} \tag{3.16}$$

Proof: See Appendix A.9.3 for the full proof.

A larger a increases the value of information relative to the bounded waiting costs and lowers equilibrium resolution probabilities. The condition $a \geq a^*$ therefore describes a regime in which resolution is sufficiently low for the proportional ordering to hold. In this regime, the appendix approximation separates the proportional effects of fewer peers and greater public precision. I show the theorem graphically in Appendix Figure 8 , which illustrates the $n = 0$ comparison across initial public precisions and levels of risk aversion. At large a, we see a strictly downward-sloping curve showing that the proportional effect is greater when initial precision is low.

The reduction in the uncertainty discount discussed in Section 3.6.3 again provides the intuition. When γ_t is low, one resolution reduces the uncertainty discount by more. To connect this gain to the proportional response, the large-a limit in Appendix A.9.3 separates the log change into a fewer-peers term, $\log((N_t - n - 1)/(N_t - n - 2))$, and an information term, $\log(1 + 2\delta/(\gamma_t + n\delta))$. At a given history n, the first term is the same across initial precisions, while the second is larger at lower γ_t. The proposition shows that this ordering holds at sufficiently large finite a.

3.7 Empirical Predictions

The empirical analysis tests the model using banks' post-COVID extension and resolution decisions in the CRE market. It focuses on across-period complementarity, which links earlier realized resolutions to later bank decisions, and on the two amplification results. These yield three predictions:

1. Across-period complementarity: early peer resolutions reduce later extensions among loans that remain unresolved.

2. Resolution informativeness: the reduction in extension is larger when realized resolutions reveal more public information.

3. Initial valuation uncertainty: the proportional increase in resolution probability is larger when initial public precision is lower.

I test these predictions by relating banks' later loan decisions to earlier resolutions by other banks in the same local collateral market. First, loan-level regressions examine whether earlier resolutions are associated with fewer subsequent extensions. To estimate their causal effect, I use a shift-share instrumental-variable design that combines banks' resolution propensities with their initial local loan shares. Second, I use state-level sale-price disclosure requirements to test whether earlier sale payoffs reduce subsequent extensions more when transaction prices are publicly available. This comparison examines whether the extension response depends on the information revealed by resolutions. Finally, I use predetermined local WFH exposure as a proxy for initial valuation uncertainty and test whether earlier resolutions lead to a larger proportional increase in subsequent resolution probability in more exposed markets. The next section describes the data and the construction of loan outcomes, measures of other banks' resolutions, and information proxies.

4 Data and Empirical Mapping

I construct a loan-quarter panel of bank-held commercial real estate (CRE) mortgages originated before 2020. The panel links county land records compiled by CoreLogic, now Cotality, to bank regulatory filings and property and business data from InfoGroup, Trepp, CompStak, and County Business Patterns. The final regression sample contains 626,839 loans across 2,089 counties and 2,900 parent banks. I track extensions, refinancing, payoffs, and defaults from 2020Q1 through 2025Q2 and use these loan histories to construct the variables for testing the model's predictions.

I exclude multifamily properties from the sample because their prices and returns move more closely with residential housing and they are less directly exposed to the work-from-home shock. The final sample includes office, retail, industrial, hotel, mixed-use, and other commercial properties. This section first describes the data sources and loan histories, then defines the regression sample and measures, and concludes with sample coverage and summary statistics.

4.1 Data Sources and Record Linkage

4.1.1 Data Sources

Commercial Property and Mortgage Data The primary source is CoreLogic's public-record mortgage data, compiled from county offices that record mortgages, deeds of trust, and property transfers. The records identify lenders and borrowers, loan amounts, and origination and maturity dates where available. They cover purchase financing, refinancing, modifications, and mortgage releases.

I link mortgage records to CoreLogic's transaction and tax-assessment data at the property level. Transaction records identify ownership transfers and provide transaction dates, disclosed sale

prices, and seller and buyer names. Assessment records provide property characteristics, assessed values, and foreclosure information. Together, these sources allow me to track properties and their mortgage histories.

I use Trepp to identify mortgages associated with commercial mortgage-backed securities (CMBS). Address matches identify properties securing an active CMBS loan when a mortgage originates. I exclude flagged mortgages from the regression sample and all associated variable construction.

I refine CoreLogic's property classifications using InfoGroup business-establishment listings and property information from CompStak. InfoGroup identifies establishments' locations and industries, allowing business activity at a property to inform its classification. Appendix B.1 describes the linkage and classification procedure.

Bank Regulatory Data I obtain bank financial information from Call Reports and FR Y-9C, institutional structure, type, and headquarters location from the National Information Center (NIC), and branch information from the Federal Deposit Insurance Corporation (FDIC). These sources support lender matching and the construction of bank financial measures. I aggregate capital ratios and CRE holdings at the ultimate parent-bank level, which also defines lender identity. These measures enter the stylized facts and the bank-coverage comparison.

Local CRE Markets and Business Activity CompStak provides transaction-level lease data for office, retail, and industrial properties in major U.S. cities. I use the data to characterize tenants' industry composition and changes in lease revenue by location and property type. Tenant composition helps measure exposure to work from home. Figure 13 in Appendix B.8 shows the CompStak valuation-property dataset and its overlap with CoreLogic.

I combine county–industry employment shares from County Business Patterns with tenant composition from CompStak and industry-level work-from-home (WFH) feasibility from Dingel and Neiman (2020). These inputs measure WFH exposure by county and property type.

To measure price dynamics, I use the RCA Commercial Property Price Indexes (CPPI), published by MSCI using Real Capital Analytics transaction data. CPPI measures monthly price changes using a repeat-sales approach. City-level indices cover major CRE markets, while regional indices cover smaller markets with fewer sales. I use these indices to classify price news as positive or negative and construct estimated loan-to-value ratios for controls.

4.1.2 Construction of the Loan-Quarter Panel

I construct the loan-quarter panel in three steps. First, I match the lender names recorded in mortgage filings to regulatory institutions and their ultimate parent banks. This crosswalk links the mortgage records to bank ownership and financial information. Appendix B.1 describes the matching procedure.

Second, I link mortgage filings into individual loan histories. The source records distinguish originations, refinancing, and modifications, but do not provide a consistent loan-level identifier across these events. I match refinancing, modification, and release records to prior loans using property identity, lien type, and the cleaned lender identifier. Modifications and refinancing by the same lender continue the existing history. Refinancing by a different lender ends that relationship and starts a new loan identifier. I link property-level default records from ownership transfers and assessments to the associated mortgages and end their histories at default.

Third, I expand the loan histories into a quarterly panel for 2020Q1–2025Q2. Each observation identifies an outstanding loan and records any extension, outside refinancing, default, or release in that quarter. The next subsection defines these events.

One data limitation is missing maturity dates. Maturity dates are missing from roughly 50% of origination records, reflecting omissions in underlying records or in CoreLogic's standardization. Reconstructing outstanding portfolios and loan performance requires filling missing contractual terms, as in Anenberg et al. (2025). I use reported terms where available and average terms by loan type and origination vintage otherwise. Later modifications or refinancing provide evidence of continued financing, and I adjust an imputed maturity when later mortgage records indicate that it is too short. Appendix B.2 describes the full procedure.

The resulting initial cohort contains 1,509,520 bank-held, non-CMBS loans outstanding in 2020Q1. Of these, 237,546 (15.74%) reach an imputed-maturity ending by 2025Q2. These administrative endings have no observed loan event in the available data, so I separate them from the extension and resolution events used in the regressions. The regression sample is a subset of this initial cohort: it contains 626,839 loans that remain outstanding at the beginning of 2022 and meet the matching, data-availability, and estimation requirements described below.

4.1.3 Loan Events and Outcomes

To construct the regression variables, I map loan events to "waiting" and "resolution" in the model. I distinguish one extension category and three resolution categories among the events I track. Extensions continue the incumbent lender's financing and correspond to waiting; resolutions end the loan. The following paragraphs explain the information content of each resolution category and its mapping to informative resolution events in the model. Appendix B.3 provides the detailed

classification rules.

Waiting: Extensions An extension is either a modification that postpones the existing loan's maturity or refinancing by the same lender with a maturity date later than that of the loan it replaces. Both continue financing by the incumbent lender beyond the previously scheduled maturity and therefore correspond to waiting in the model.

1. Modification with a maturity extension. Figure 12 in the appendix shows an example from Cook County land records. I identify these events using recorded modifications in the mortgage data, requiring that the modification specify a later maturity date.

2. Same-lender refinancing. The Office of the Comptroller of the Currency notes in its guidance that commercial-loan refinancing can constitute a modification or accommodation for a borrower experiencing financial difficulty. Such refinancing continues the financing of the same property, and the borrower's financial difficulties may persist after the refinancing.

Appendix B.3 describes how extensions are identified.

Resolution: Default Default resolutions are loan endings identified through foreclosure and distressed ownership transfers. Default resolutions reveal information at several stages. The ownership transfer makes the distressed outcome observable to other lenders whose collateral is exposed to similar local conditions. For some transfers, I observe sales of the real estate owned (REO) properties through private transactions or auctions one to three years later. These sales provide additional information about collateral value, particularly when purchase prices are disclosed. We can interpret the information content of the default process as follows: the distressed transfer reveals the outcome of a lending relationship, while a subsequent sale provides more direct information about asset value.

I record the date of each default event as the date when ownership passes from the borrower to the lender. Since this date marks the transition from lending to recovering value through ownership of the collateral, it marks the starting point of the resolution of the lending relationship and makes the distressed outcome observable. This also provides a consistent timing measure regardless of whether a subsequent REO sale is observed.

Resolution: Sale Payoff I define a loan end as a sale payoff when the existing loan is repaid through refinancing by a different lender or a recorded mortgage release , and the property is sold in the same quarter. The definition combines evidence of loan termination with a property sale observed within the same quarter. This event marks the end of the previous lending relationship with the bank, and at the same time a disclosed transaction price provides a valuation benchmark for comparable properties. This provides a benchmark for other lenders to reassess their collateral

and subsequent extension or resolution decisions.

Other Resolutions *Other resolutions* comprise repayment through refinancing by a different lender or a recorded mortgage release without a property sale in the same quarter. These events do end the existing loan and the lender's exposure, but their information content is more limited compared to the previous two types of resolution events. For example, external refinancing may reveal information through the new loan terms, but it does not reveal a sale price for the property, so information from these events can be much noisier. In the variable construction, I include only defaults and sale payoffs in the main measure of earlier informative resolutions.

Additionally, another potential "resolution event" is the loan reaching maturity. I allow each loan to leave the panel at maturity when no subsequent extension, refinancing, default, or other termination event is observed. I also exclude these exits from resolution counts because maturity does not clearly indicate the end of the lending relationship, unlike an observed refinancing, mortgage release, or default. These endings may reflect missing observations. For example, the absence of a later filing may reflect an unobserved extension or refinancing.

4.2 Regression Panel and Measures

4.2.1 Sample and Timing

The empirical predictions from the model connect earlier resolutions to subsequent lending decisions: the tests ask whether more earlier resolutions predict less extension, whether the relationship is stronger when resolutions reveal more public information, and whether the response varies with initial uncertainty. To define the timing for "earlier resolution" and "subsequent lending", I separate information arrival from subsequent decisions using two fixed observation windows.

More precisely, I follow bank-held CRE loans outstanding at the beginning of 2020 through 2025Q2. The earlier window, t_1, covers 2020Q1–2021Q4; the later window, t_2, covers 2022Q1–2025Q2. Since resolutions can occur repeatedly within a market, a quarterly analysis makes it difficult to separate the effects of information arriving at different times (see the discussion in Callaway and Sant'Anna (2021)). I therefore use a fixed-window design, counting the total number of resolutions during t_1 and relating it to loan outcomes during t_2. This keeps the exposure measure fixed before the outcome window and allows me to test whether information from earlier resolutions affects subsequent decisions for loans that remain outstanding. In other words, we are measuring the extension probability conditional on survival in t_1, which corresponds to the proposition target q^* that supports the predictions.

One potential concern with the cutoff is that the later window largely overlaps with the monetary-tightening cycle. Higher rates can make refinancing harder and increase incentives to extend loans.

This concern also arises in a dynamic analysis because much of the post-WFH sample overlaps with the rate hikes. I therefore address this concern through the empirical specifications, distinguishing common changes in extension activity from differences in exposure to tightening that may be related to earlier resolutions. Section 5 discusses the roles of fixed effects, local controls, and the instrumental-variable design in this comparison. I also use alternative cutoffs between t_1 and t_2 to check whether the results depend on the chosen cutoff.

4.2.2 Variable Definitions

The regression variables follow this timing: the outcome records each loan's extension during t_2, and earlier exposure counts informative resolutions during t_1 in the same market.

Extension and Earlier Resolutions For each loan i outstanding at the beginning of 2022, $Ext_i = 1$ if the incumbent lender extends its maturity at least once during t_2. Otherwise, $Ext_i = 0$. The indicator records whether financing continues beyond the previous maturity at any point in the later window.

I define a local CRE market as a county-property type pair, denoted as $(c(ounty), p(type))$. Properties within a market share exposure to similar local demand and conditions, so learning is relevant to lenders that hold collateral in the same market. For a loan held by parent bank b, I count defaults and sale payoffs during t_1 on loans held by other parent banks in that market. The measure $SaleCounts_{c,p,-b,t_1}$ counts sale-payoff events, which are loan repayment events with a same-quarter property transaction; $DefaultCounts_{c,p,-b,t_1}$ counts defaults. The total earlier-resolution count is

$$ResCounts_{c,p,-b,t_1} = SaleCounts_{c,p,-b,t_1} + DefaultCounts_{c,p,-b,t_1}.$$

The subscript $-b$ indicates that I exclude loans held by the same parent bank. That is, for each loan extension outcome, I count only the previous period's resolution events from its peer banks, aiming to limit the mechanical influence of the loan's own parent bank's resolutions on its measured exposure. The regression panel is at the loan level. Each loan is matched to the resolution count for its county and property type and related to its subsequent extension outcome indicator, along with the controls.

Loan and Market Controls Loan controls account for differences in the composition of outstanding loans across banks and markets. More precisely, I control for the loan age, loan-to-value (LTV) at origination, and time to maturity. These measures capture the loan life-cycle, initial leverage, and proximity to contractual repayment. Origination LTV compares the original loan amount with the property's value at origination. Appendix B.5 describes its construction.

At the county–property-type level, I control for property-market liquidity, bank concentration, and the initial stock of other-bank loans. Liquidity is measured by annual turnover and log transaction counts from 2015-2019, which captures market activity and the ease of selling collateral. Bank concentration is measured by the Herfindahl–Hirschman index of banks' shares of total CRE loan amounts in the market. The initial other-bank loan stock, $N_{c,p,-b,t_0}$, is measured at the beginning of 2020.

Another important control is the predetermined exposure to the WFH shock. A market is more exposed when its pre-pandemic demand for space comes from industries whose jobs can be performed from home. I therefore combine county employment composition, industries' use of different property types, and their WFH feasibility using pre-pandemic weights.

Let $s_{c,j}$ denote industry j's share of county c's employment in 2019, $u_{j,p}$ its lease-value share for property type p in CompStak during 2009–2019, and h_j its WFH feasibility from Dingel and Neiman (2020). Exposure is the weighted average

$$H_{c,p} = \frac{\sum_j s_{c,j} u_{j,p} h_j}{\sum_j s_{c,j} u_{j,p}}. \tag{4.1}$$

Exposure is higher when industries with greater WFH feasibility account for a larger share of local demand for that property type. For example, an office market is more exposed when local employment is concentrated in industries that use office space and can readily work from home.

The roles of the controls are discussed in detail in the next section. Additionally, the regression specifications require an initial other-bank loan stock of at least 10, mainly due to the IV construction process. Appendix B.4 documents the construction and sample restrictions; Section 5 describes the specifications and reports their observation counts.

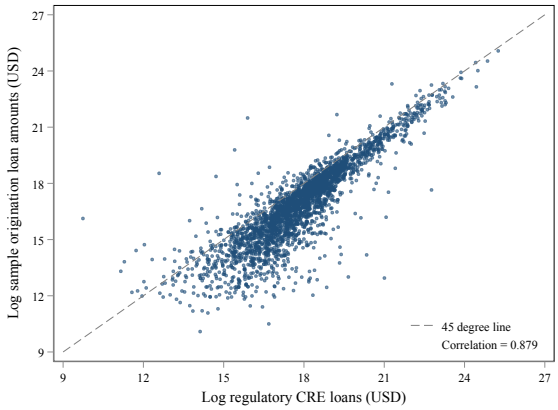
4.3 Coverage and Summary Statistics

I assess bank coverage by comparing sampled loan amounts with regulatory CRE holdings. I aggregate origination amounts at the parent-bank level and match them to nonfarm nonresidential CRE balances from Call Reports and FR Y-9C in 2020Q1. Panel A of Figure 6 plots the two measures in log. Their correlation is 0.879: banks with larger regulatory CRE portfolios also tend to have larger total origination amounts in the sample. We also see that the Corelogic implied loadings compared to the regulatory holdings are lower for some small and regional banks, which are reasonable as those banks can have more holdings on the small assets, which can be missed from the public record, but the overall coverage of the Corelogic is decent and reasonable.

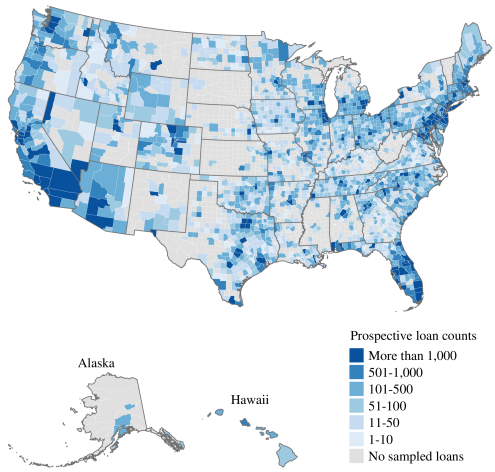
The regression sample contains 626,839 loans across 2,089 counties. Panel B of Figure 6 shows their geographic coverage. Loans are concentrated in major metropolitan areas, coastal regions, and Sun Belt markets, and the sample also covers less populated regions. In the regression sample, the

Figure 6: Coverage of the CRE loan sample

Panel A: Bank coverage relative to regulatory CRE holdings



Panel B: Geographic distribution of CRE loans



Notes. *Panel A.* Each point represents a parent bank. The vertical axis shows the sum of origination amounts for loans in the regression sample, grouped by the sample's parent-bank identifier. The horizontal axis shows the (parent) banks' regulatory nonfarm nonresidential CRE loan balance in 2020Q1. Both axes use natural logarithms of dollar amounts. The line is the 45 degree line; the reported correlation is calculated from the plotted logarithms. Appendix B.7 describes the construction.

Panel B. Counties are shaded by the number of loans. This prospective coverage map contains the 626,839 regression-sample loans across 2,089 counties.

average county contains roughly 300 loans and the median contains 83.

Table 1 summarizes the 626,839 loans in the regression sample. During 2022Q1–2025Q2, 3.3% experienced at least one extension. The mean origination amount is $1.428 million, compared with a median of $0.420 million, reflecting a right-skewed distribution. Office and retail properties account for the largest shares, followed by mixed use and other commercial properties at 18.8% and industrial properties at 14.9%.

Panel A also describes the information available from other banks' resolutions. The average loan faces an initial stock of 1,721 other-bank loans in its county–property-type market. Before the extension window, these peers generate an average of 65 informative resolutions, comprising 50 sale payoffs and 15 defaults.

Panel B of Table 1 describes the 2,900 represented parent banks and their regulatory categories. Assets are observed for 2,786 banks, with a median of $312 million. The sample includes substantial representation of community and regional banks. CRE lending is an important part of these banks' portfolios: the average regulatory CRE-to-asset ratio is 18.36%.

Table 1: Summary statistics

Variable	Unit	Mean	SD	Median	N
Panel A. Loan-level summary statistics					
Loan outcomes and characteristics:					
Loan extension	$\{0,1\}$	0.03	0.18		626,839
Origination amount	USD mn	1.43	10.54	0.42	626,839
Origination LTV	%	75.75	596.38	59.21	626,839
Loan age	Years	6.95	4.43	6.00	626,839
Minimum remaining maturity	Years	3.93	4.83	2.50	626,839
Property type:					
Office		0.24			626,839
Retail		0.40			626,839
Industrial		0.15			626,839
Hotel		0.03			626,839
Mixed use and other commercial		0.19			626,839
Earlier resolutions and initial exposure:					
Other-bank informative resolutions	Count	65.05	112.98	27.00	626,839
Other-bank sale payoffs	Count	50.33	78.27	21.00	626,839
Other-bank defaults	Count	14.72	39.72	5.00	626,839
Initial other-bank loan stock	Count	1,721.88	3,456.22	626.00	626,839
Local market characteristics:					
WFH exposure	0-1	0.23	0.10	0.23	626,839
Pre-pandemic annual property turnover	%	2.36	1.95	2.26	626,839
Average annual purchases, 2015-2019	Count	123.39	254.90	35.60	626,839
Bank-lending HHI	0-1	0.13	0.12	0.09	626,839
Panel B. Represented parent banks (2019Q4 financials)					
Total assets	USD bn	7.64	93.96	0.31	2,786
Regulatory CRE loans/assets	%	18.36	11.63	16.97	2,786
CET1 capital buffer	pp	10.28	8.02	7.98	2,782
Counties with sampled loans	Count	19.23	58.53	6.00	2,900
Bank regulatory type					2,900
Large banks		0.01			40
Regional and large community banks		0.13			377
Small community banks		0.86			2,483
Panel C. Represented counties (2019 CBP)					
Employment	Thousands	58.04	178.05	11.97	2,089
Establishments	Thousands	3.58	10.61	0.89	2,089
Sampled loans	Count	300.07	783.84	83.00	2,089

Notes: Observations are equally weighted within each panel. Extensions cover 2022Q1–2025Q2; earlier resolutions cover 2020Q1–2021Q4 for other banks in the same county and property type. Loan age uses the retained later-window observation; remaining maturity is the minimum across retained observations in that window. Annual turnover is average annual purchases in 2015–2019 divided by the property stock.
The CET1 capital buffer is the CET1 ratio minus the bank-specific requirement, including the conservation buffer and G-SIB surcharge. Bank types use available 2015Q1–2025Q3 asset histories: large banks have maximum assets of at least USD 100 billion; among the remainder, regional and large community banks have mean assets above USD 1.284 billion (the 2019 CRA threshold), and small community banks have mean assets at or below it.

5 Empirical Design and Results

This section tests the empirical predictions of the model in Section 3. I first examine the relationship between earlier resolutions at other banks and subsequent loan extension using OLS and a bank-based instrumental-variables design. I then test whether the extension response is stronger where transaction prices are publicly disclosed and whether the proportional resolution response is larger in markets with greater initial valuation uncertainty. The results are consistent with the model's predictions. Finally, I compare sale payoffs with good and bad price news, defined using transaction prices relative to adjusted assessed values. This comparison tests whether adverse valuation news alone can explain the findings.

5.1 Prediction 1: Earlier Peer Resolutions and Subsequent Extension

Prediction 1: Early peer resolutions reduce later extension among loans that remain unresolved.

This prediction follows from across-period complementarity. Holding the initial state variables N, γ and the other model parameters fixed, random realizations of waiting costs c_i generate differences in early peer resolutions. Intuitively, more early resolutions reveal more public information, lowering the benefit of waiting and thus reducing the extension rate in the next period.

To test this prediction, I estimate the following specification:

$$\begin{aligned} Ext_{i,c,p,b,t_2} = \alpha + \beta ResCounts_{c,p,-b,t_1} + \Theta' W_{i,c,p,b,t_0} \\ + \lambda_c + \kappa_p + \psi_b + \varepsilon_{i,c,p,b,t_2}, \end{aligned} \tag{1}$$

Here, i indexes loans, c counties, p property types, and b parent banks. The sample, timing, and variables follow Section 4.2. The earlier window, t_1, covers 2020Q1–2021Q4, and the later window, t_2, covers 2022Q1–2025Q2. The sample consists of loans originated before the pandemic that remain outstanding at the beginning of 2022Q1. A market is a county–property-type pair.

The dependent variable in the regression is Ext_i, which records whether a loan is extended during t_2. It equals one if the loan is extended through a modification or same-parent refinancing during t_2, and zero otherwise. The main explanatory variable is $ResCounts_{c,p,-b,t_1}$. It counts sale payoffs and defaults at other banks in the same market during t_1. The $-b$ shows that this counts excludes the resolution events from the loan's own parent bank, which aims to limit the mechanical influence of its own resolutions on measured peer exposure.

The coefficient of our interest is β. Prediction 1 implies $\beta < 0$: conditional on the loan being outstanding at the beginning of t_2, more peer resolutions in t_1 are associated with a lower probability of subsequent extension. The prediction holds conditional on the aggregate state variables, including the initial loan stock N and public uncertainty γ. I therefore control for initial market size and

predetermined WFH exposure. The latter combines local tenant-industry composition with industry exposure to the WFH shock, as described in Section 4.2.

Besides the controls for size and WFH exposure, the control vector W also includes loan-level and market-level characteristics. The loan-level controls include loan age, loan-to-value (LTV) at origination, and time to maturity. The market-level controls includes measures of property market liquidity and bank concentration. The specification also includes county, property-type, and parent-bank fixed effects. County fixed effects absorb persistent differences in local economic conditions and lending practices. The coefficient β is therefore identified from variation within counties: property types differ in their earlier resolution counts, and banks face different peer counts because each bank is excluded from its own exposure measure. property type fixed effects absorb unobservables within each property type, while parent-bank fixed effects account for common unobservables at the bank level. Standard errors are clustered by county to allow loan outcomes within local property markets to be correlated.

The main endogeneity concern is that local economic conditions may affect both earlier resolutions and later extensions. For example, a persistent decline in property liquidity may affect how many loans banks resolve during t_1 and whether they extend the remaining loans during t_2. Although I include local controls to capture market characteristics that may create this concern, some unobserved market conditions may still affect both decisions. I therefore interpret the OLS estimates as correlations rather than causal effects.

5.1.1 Instrument Variable Construction

To address the potential endogeneity of earlier resolutions, I construct instruments using market level exposure heterogeneity on the bank level shock. The motivation of exogenous shock comes from model setup: we model the randomness in the realization of resolution from banks' i.i.d distributed c, such that banks with higher waiting costs resolve loans more readily. Map to the empirics, we want to find the bank level shock that impact the resolution probability in the previous period. Since we have banks that span multiple markets, banks that resolve more loans elsewhere may also resolve more loans in the local market. Combining these outside-market propensities with banks' initial local loan holdings yields predicted resolution activity at market level.

More precisely, I construct a shift-share instrument follows Greenstone et al. (2020). The shifter is the bank-level resolution propensity shock, estimated by separating bank and market effects in resolution rates. Let $m = (c, p)$ denote the county–property-type market and k index banks. Excluding market m, estimate across markets $\ell \neq m$:

$$\frac{ResCounts_{\ell,k,t_1}}{N_{\ell,k,t_0}} = d_\ell + Shock_k + e_{\ell,k}, \qquad \ell \neq m,$$

where N_{ℓ,k,t_0} is bank k's initial loan stock in market ℓ. Market fixed effects d_ℓ absorb differences in resolution rates common to banks operating in the same market. The estimated bank-level fixed-effect, $\widehat{Shock}_k^{-m}$, captures bank k's propensity to resolve loans relative to other banks, using its lending activity outside market m.

Next, I construct initial bank shares in each market, excluding the loan's own parent bank:

$$Share_{m,k,t_0}^{-b} = \frac{N_{m,k,t_0}}{N_{m,-b,t_0}}, \qquad \text{where } N_{m,-b,t_0} = \sum_{k \neq b} N_{m,k,t_0}.$$

I then aggregate shifter and the shares for each market-bank pair. The instruments for both the resolution rate and resolution number

$$Z_{m,-b}^{R} = \sum_{k \neq b} Share_{m,k,t_0}^{-b} \widehat{Shock}_k^{-m}, \qquad Z_{m,-b}^{X} = N_{m,-b,t_0} Z_{m,-b}^{R}.$$

We include the rate instrument as the bank level shock is better estimated through the resolution propensity instead of count. The rate instrument is larger when banks with higher outside-market resolution propensities initially hold a larger share of local loans. Multiplying by the initial loan stock gives the corresponding count instrument. The count instrument captures variation in information quantity, the key object in this test.

In the full specification, I jointly instrument for earlier resolution counts and the resolution rate, including both instruments in each first-stage regression. The peer resolution rate is $R_{m,-b,t_1} = ResCounts_{m,-b,t_1} / N_{m,-b,t_0}$, reported in percentage points. Both the count and rate enter the full specification. The market size control are used for separate the impact market size on the counts. I include the same controls and fixed effects as in the OLS regressions. The standard errors are clustered at county level to align with the OLS result [2] .

5.1.2 Instrument Validity

First Stage Because the specification includes both resolution counts and the resolution rate, identification requires that the two instruments generate sufficiently distinct variation in these endogenous variables, conditional on the controls and fixed effects. Both instruments enter each first-stage regression.

Table 2 reports the first-stage estimates. The rate instrument positively predicts the resolution rate, with a coefficient of 0.444 (standard error 0.110), while the count instrument positively predicts

[2]The clustering / inference for shift-share instrument worth caution, as two geographically distant markets can have correlated instruments because they are exposed to the same banks. County clustering allows dependence within each county, but does not capture this connection across counties. To mitigate that concern, I also implement the clustering proposed by Adão, Kolesár, and Morales (2019) using the `ivreg_ss` Stata module (Adão, Kolesár, Morales, & Zhang, 2020) as the robustness test, and show that the coefficient are still significant at 1%.

Table 2: Joint Count–Rate First Stages and Diagnostics

	Rate R (pp)	Count X (SD)
Rate instrument Z^R	0.44366^{***}	-0.00515
	(0.10986)	(0.00442)
Count instrument Z^X	-0.03188	0.38786^{***}
	(0.02093)	(0.10402)
Excluded-instrument F	9.46	7.38
Conditional SW F	16.25	25.14
Observations	626,839	626,839
County clusters	2,089	2,089
Joint diagnostics	Statistic	p-value
KP rank LM	11.65	0.0006
KP Wald F	8.12	—
Anderson–Rubin F	4.06	0.0174
Stock–Wright LM	5.90	0.0522

Notes. Both instruments enter both first stages, together with WFH exposure, loan and local controls, initial peer-stock controls, and county, property-type, and parent-bank fixed effects. The rate instrument is multiplied by 100; the count instrument is divided by $s_I = 112.9793$. Standard errors are clustered by county. The Anderson–Rubin and Stock–Wright tests concern the joint null for both endogenous coefficients, conditional on valid instruments. $^{**}p < 0.05$, $^{***}p < 0.01$.

resolution counts, with a coefficient of 0.388 (standard error 0.104). The cross-instrument coefficients are small and statistically insignificant. The first stage coefficients suggest that the instruments capture distinct variation in the resolution rate and the number of realized resolutions.

I report several first-stage statistics and tests for the joint 2SLS specification. The Sanderson–Windmeijer conditional first-stage F-statistics are 16.25 for the rate and 25.14 for the count, both above the conventional rule-of-thumb threshold of 10. The Kleibergen–Paap rank LM test rejects underidentification ($p = 0.0006$). However, the Kleibergen–Paap Wald F-statistic is 8.12, so weak-instrument concerns remain. I therefore also report the Anderson–Rubin test, which rejects the joint null that the coefficients on resolution counts and the resolution rate are both zero ($p = 0.0174$). This provides evidence against the joint null using inference that is robust to weak instruments.

Exclusion Restriction The exclusion restriction requires that, conditional on the controls and fixed effects, exposure to other banks' outside-market resolution propensities affects subsequent loan extension only through earlier local resolution counts and rates. Under an exposure-based interpretation of the shift-share design (Goldsmith-Pinkham, Sorkin, & Swift, 2020), identification relies on differences in initial local bank shares being unrelated to unobserved determinants of subsequent extension, after conditioning on these covariates.

The instrument construction limits two sources of mechanical dependence. When constructing the instrument for market m, I exclude that market from the estimation of bank resolution propensities. Local resolution outcomes therefore do not directly enter the estimated shifters. I also

exclude the loan's parent bank when aggregating exposure, so its own resolution propensity does not directly enter the instrument.

One potential violation arises if banks with greater resolution pressure also reduce refinancing supply. Exposure to these banks could affect subsequent extension through borrowers' outside financing options, independently of earlier resolutions. If reduced refinancing availability induces incumbent lenders to extend existing loans, this channel predicts greater extension, opposite to the negative count coefficient predicted by the information mechanism. This argument limits the ability of that particular channel to explain a negative estimate.

A second concern is bank sorting across markets. Banks with higher resolution propensities may initially hold larger shares in markets facing different subsequent economic conditions. I control for predetermined WFH exposure, pre-pandemic property-market liquidity, bank concentration, and initial loan stock, together with county, property-type, and parent-bank fixed effects. These controls account for observable differences across markets and average differences across locations, sectors, and lenders. The remaining identifying assumption is that initial exposure is uncorrelated with unobserved determinants of subsequent extension. An additional robustness exercise would reconstruct the instruments using earlier bank shares to assess sensitivity to portfolio allocation immediately before the pandemic.

5.1.3 Results

Table 3 reports the relationship between earlier resolutions at other banks and subsequent loan extension. All specifications use the same sample of 626,839 loans, whose summary statistics are reported in Section 4.3. All columns include county and property-type fixed effects, loan characteristics, initial loan-stock controls, and WFH exposure. Column 2 adds parent-bank fixed effects. Column 3 further includes pre-pandemic property-market liquidity, bank concentration, and the earlier resolution rate. Column 4 jointly instruments for resolution counts and the resolution rate. The extension indicator is multiplied by 100, so coefficients measure percentage-point changes in extension probability. Resolution counts and WFH exposure are scaled by their respective standard deviations.

Earlier resolutions are negatively associated with subsequent extension in every specification. In column 1, a one-standard-deviation increase in earlier resolution counts is associated with a 0.278 percentage-point reduction in extension probability. Adding parent-bank fixed effects leaves the coefficient nearly unchanged at -0.272 and reduces its standard error from 0.138 to 0.077. With local controls and the resolution rate included, the coefficient is -0.264, significant at the 1 percent level.

The full-control estimate implies a 0.264 percentage-point reduction in extension probability for approximately 113 additional earlier resolutions. This is just under one-tenth of the mean extension rate of 3.05 percent. The magnitude is comparable to the coefficient on WFH exposure: across the

Table 3: Earlier Peer Resolutions and Subsequent Extension

	(1) OLS	(2) OLS	(3) OLS Full controls	(4) Joint IV Full controls
Resolution counts (SD)	-0.278^{**}	-0.272^{***}	-0.264^{***}	-0.229^{***}
	(0.138)	(0.077)	(0.076)	(0.066)
WFH (SD)	0.257^{*}	0.314^{**}	0.349^{**}	0.505^{***}
	(0.144)	(0.137)	(0.141)	(0.178)
Resolution rate R (pp)			-0.025^{**}	-0.306
			(0.012)	(0.192)
Parent-bank fixed effects	No	Yes	Yes	Yes
Local controls	No	No	Yes	Yes
Initial peer-stock controls	Yes	Yes	Yes	Yes
Peer resolution rate	No	No	Yes	Yes
Observations	626,839	626,839	626,839	626,839
R^2	0.0451	0.0729	0.0729	
KP Wald F				8.12
Conditional SW F: count				25.14
Conditional SW F: rate				16.25

Notes. Outcome: extension during 2022Q1–2025Q2. Coefficients and standard errors are in percentage points. Earlier peer events occur during 2020Q1–2021Q4 and exclude the loan's own parent bank. All count regressors use the common total-event SD of 112.9793; rates are in percentage points. All columns include county and property-type fixed effects, log loan age, origination LTV, raw minimum quarters to maturity, and standardized WFH exposure. Local controls are pre-pandemic turnover, log transaction volume, and bank concentration. All columns include initial peer-stock controls. Column 4 instruments both the count and rate, using both instruments in both first stages. Standard errors clustered by county are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

OLS specifications, a one-standard-deviation increase in exposure is associated with a 0.257–0.349 percentage-point increase in extension probability.

The IV estimate is also negative and similar in magnitude to the OLS estimate. Column 4 reports a count coefficient of −0.229 (standard error 0.066), compared with −0.264 in the corresponding OLS specification. Conditional on the resolution rate and the remaining covariates, a one-standard-deviation increase in earlier resolution counts corresponds to a 0.229 percentage-point reduction in subsequent extension, between one-twentieth and one-tenth of the mean extension rate. Thus, the negative relationship also appears when resolution activity is instrumented using other banks' outside-market resolution propensities.

The resolution-rate coefficient is negative in both the OLS and IV specifications, although the IV estimate is larger in absolute magnitude. Overall, the negative count coefficients support the first prediction: more earlier resolutions at other banks are followed by a lower probability of subsequent loan extension. This pattern is consistent with the model's across-period complementarity, whereby realized resolutions reduce the incentive of remaining lenders to delay resolution. Appendix D.1 reports sensitivity to alternative boundaries between the earlier and later windows.

5.2 Prediction 2: Transaction-Price Disclosure

Prediction 2: The reduction in extension is larger when realized resolutions reveal more public information.

The second prediction concerns the informativeness of earlier resolutions. If information revealed by resolutions contributes to the decline in subsequent extensions, the relationship should be stronger when resolution results are more informative concerning the collateral value. I examine this prediction using differences in transaction-price disclosure policies across states.

More precisely, whether real estate transaction prices are publicly available varies across states. In disclosure states, sale prices are generally available through public property or transfer records, allowing lenders and other market participants to observe prices paid for comparable properties. In nondisclosure states, an ownership transfer may be recorded without making the transaction price publicly available. The relevant distinction is the information available to lenders outside the transaction: buyers, sellers, and their lenders can observe the agreed price under either regime, but its availability to other lenders differs (Gárate, Hilterbrand, & Pennington-Cross, 2025; Ling, Naranjo, & Scheick, 2021).

Differences in disclosure reflect state recording, taxation, and confidentiality rules. Transaction prices may appear in recorded deeds or accompanying transfer documents, or be inferred from public transfer taxes based on transaction value. Public price availability therefore does not require a separately reported sale-price field. Disclosure policies can reflect the competing goals of accurate property assessment, transparency, and taxpayer privacy, which are generally unrelated to other CRE-market characteristics such as liquidity and bank market power. Access to transaction prices can help assessors value properties and owners compare their assessments with similar sales (Berrens & McKee, 2004) . Some states restrict public access even when prices are reported to tax authorities. A Utah legislative audit discusses this trade-off between valuation transparency and property owners' privacy (Office of the Legislative Auditor General, 2024, p. 21) . I classify Alaska, Idaho, Kansas, Louisiana, Mississippi, Montana, New Mexico, Texas, Utah, and Wyoming as nondisclosure states. This group includes states where transaction prices are not generally required to be reported to public authorities, as well as states where prices are reported to some public authorities but public access to that information is restricted (Kansas, Montana, New Mexico, and Wyoming). Besides those states, another complicated case is Missouri. Missouri has a mixed system, with disclosure requirements and price availability varying across local jurisdictions (Gárate et al., 2025, Section 2). The baseline state-level classification treats Missouri as a disclosure state, reflecting public price availability in its several major markets, including St. Louis County, St. Charles County, Jackson County, and the City of St. Louis.

I use differences in transaction-price disclosure policies to test whether earlier sale payoffs are more strongly associated with subsequent extension where prices are publicly available. Defaults

do not themselves reveal transaction prices, so the information they provide is not directly affected by price-disclosure policies.

To test this prediction, I retain the same loan sample as in Prediction 1, replace total earlier peer-resolution counts with sale-payoff counts, and interact the latter with price-disclosure status:

$$\begin{aligned} Ext_{i,c,p,b,t_2} = \alpha &+ \beta_S SaleCounts_{c,p,-b,t_1} \\ &+ \theta_S \left(SaleCounts_{c,p,-b,t_1} \times D_c \right) + \Theta' W_{i,c,p,b,t_0} \\ &+ \lambda_c + \kappa_p + \psi_b + \varepsilon_{i,c,p,b,t_2}, \end{aligned}$$

Here, $SaleCounts_{c,p,-b,t_1}$ counts earlier sale payoffs at other banks in the same county–property-type market, and D_c indicates counties classified as subject to public sale-price disclosure. The specification uses the same loan and market controls and fixed effects as the prediction 1. Standard errors are clustered by county.

The coefficient of interest, θ_S, measures the difference in the relationship between earlier sale payoffs and subsequent extensions across disclosure and nondisclosure markets. The coefficient on earlier sale payoffs is β_S in nondisclosure markets and $\beta_S + \theta_S$ in disclosure markets. Prediction 2 implies $\theta_S < 0$, that is, an additional sale payoff is associated with a larger reduction in subsequent extensions where its transaction price is publicly available.

5.2.1 Results

Table 4 reports the results for Prediction 2. More precisely, it shows the relationship between earlier sale payoffs and subsequent extensions across disclosure regimes. All specifications use the same loan sample as Prediction 1 and include county and property-type fixed effects, loan characteristics, initial loan-stock controls, and WFH exposure. Column 2 adds parent-bank fixed effects, and column 3 adds local controls and the resolution rate at other banks. Sale-payoff counts are divided by the same total-event standard deviation as in Prediction 1, so the scaling remains the same. The coefficients are reported in percentage points and indicate the change in extension probability associated with an increase of around 113 sale-payoff events.

Across all three specifications, the sale-payoff–disclosure interaction is negative and statistically significant at the 5 percent level. The estimate is −0.424 in the baseline in column 1. Its magnitude decreases slightly to 0.389 after adding parent-bank fixed effects and increases to 0.418 after including the full controls. The similar magnitude across specifications indicates a consistently more negative relationship between earlier sale payoffs and subsequent extensions in disclosure markets, where each sale-payoff resolution event carries more precise information on collateral value to the public.

As for the magnitude, approximately 113 additional sale payoffs at other banks are associated

Table 4: Transaction-Price Disclosure and Peer Sale Payoffs

	(1) Baseline	(2) + Bank FE	(3) + Local controls and rate
Sale payoffs × disclosure	−0.424**	−0.389**	−0.418**
	(0.205)	(0.184)	(0.188)
Sale payoffs (SD)	−0.586***	−0.569***	−0.493***
	(0.154)	(0.136)	(0.148)
WFH (SD)	0.315**	0.351***	0.405***
	(0.138)	(0.132)	(0.140)
Resolution rate R (pp)			−0.017
			(0.012)
Parent-bank fixed effects	No	Yes	Yes
Local controls	No	No	Yes
Initial peer-stock controls	Yes	Yes	Yes
Peer resolution rate	No	No	Yes
Observations	626,839	626,839	626,839
R^2	0.0451	0.0728	0.0728

Notes. Outcome: extension during 2022Q1–2025Q2. Coefficients and standard errors are in percentage points. Earlier peer events occur during 2020Q1–2021Q4 and exclude the loan's own parent bank. All count regressors use the common total-event SD of 112.9793; rates are in percentage points. The disclosure interaction compares the sale-count slope across disclosure regimes. The disclosure main effect is absorbed by county fixed effects. All columns include county and property-type fixed effects, log loan age, origination LTV, raw minimum quarters to maturity, and standardized WFH exposure. Local controls are pre-pandemic turnover, log transaction volume, and bank concentration. All columns include initial peer-stock controls; column 3 also includes the peer resolution rate. Standard errors clustered by county are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

with a 0.493 percentage-point reduction in extension probability in nondisclosure markets in the full specification. The corresponding reduction in disclosure markets is 0.911 percentage points, which is the sum of the sale-payoff coefficient and its disclosure interaction. This is almost twice the magnitude in nondisclosure markets, suggesting that publicly available prices can explain a nontrivial part of the correlation between extensions and earlier resolutions.

These findings support the second prediction: earlier sale payoffs are more negatively associated with subsequent extension where transaction prices are publicly available. One concern is that systematic differences between disclosure and nondisclosure markets, including shocks specific to Texas, may contribute to this pattern. As an additional comparison, Appendix Table 9 jointly includes sale payoffs, defaults, and their interactions with disclosure status. Unlike sale payoffs, borrower-to-creditor transfers do not themselves reveal an arm's-length market price. The sale-payoff interaction remains negative and significant, whereas the default interaction is positive and statistically insignificant in the full specification. Thus, the stronger negative relationship in disclosure markets appears for sales that can reveal transaction prices, but not for defaults, consistent with the public-information mechanism.

5.3 Prediction 3: Initial Valuation Uncertainty

Prediction 3: The proportional increase in resolution probability is larger when initial public precision is lower.

The third prediction concerns how initial valuation uncertainty affects lenders' responses to earlier resolutions. When the economy starts with higher uncertainty, information revealed by other banks' resolutions has a greater effect on the incentive to resolve. This can be interpreted as a higher marginal value of additional information when uncertainty is high. Note that this prediction concerns proportional changes in resolution probability, rather than the level changes in the previous two predictions. The distinction matters because when the initial resolution probability is low, a larger proportional response does not necessarily imply a larger increase in levels. I therefore examine whether initial uncertainty amplifies the response relative to the underlying resolution probability.

I use predetermined WFH exposure as a proxy for differences in initial valuation uncertainty across markets. The shift toward remote work is a major source of valuation uncertainty for commercial real estate loans. Higher exposure can lower demand for commercial space, reducing rents and occupancy. This can increase the need for conversion or renovation of local properties, creating more uncertainty about their future values.

I construct the exposure measure using pre-pandemic information to avoid look-ahead bias. It combines pre-pandemic industry composition with industry-level WFH feasibility and does not depend on subsequent property-market outcomes.

The specification to test the prediction uses the interaction design follows the Rajan and Zingales (1998). More precisely, I test whether earlier resolutions have a stronger association with subsequent resolutions in markets where the mechanism predicts a greater role for information. I interact earlier resolution counts with predetermined WFH exposure. An OLS regression, which is a linear probability model with this interaction, would test differences in responses in levels. To test the prediction on proportional responses, I specify an exponential conditional mean and estimate it using Poisson pseudo-maximum likelihood (PPML). The interaction coefficient measures how the proportional response to earlier resolutions varies with WFH exposure:

$$\mathbb{E}\left[Resolution_{i,c,p,b,t_2} \mid \mathcal{X}_i\right] = \exp\left\{\begin{array}{l} \alpha + \zeta_H\left(ResCounts_{c,p,-b,t_1} \times WFH_{c,p}\right) \\ + \zeta_X ResCounts_{c,p,-b,t_1} + \rho WFH_{c,p} + \Theta' W_{i,c,p,b,t_0} \\ + \lambda_c + \kappa_p + \psi_b \end{array}\right\}.$$

The $Resolution_{i,c,p,b,t_2}$ dummy is defined in a similar way to the extension dummy introduced earlier. It equals one if the loan experiences a default or sale payoff during the later period, and

zero otherwise. I use the same event definition as for earlier informative resolutions to keep the definitions consistent. The specification also controls for the resolution count and WFH exposure separately, along with the same loan and market controls and county, property-type, and parent-bank fixed effects. Standard errors are clustered by county.

The coefficient of interest is ζ_H. Under this specification,

$$\frac{\partial \log \mathbb{E}[Resolution_i \mid \mathcal{X}_i]}{\partial ResCounts} = \zeta_X + \zeta_H WFH_{c,p}.$$

The model prediction leads to a positive ζ_H. That is, the earlier resolutions are associated with a larger proportional increase in subsequent resolution probability when the local market is more exposed to the WFH shock.

5.3.1 Results

Table 5 reports weighted PPML estimates of subsequent resolution probability, defined as a default or sale payoff during the later period. All columns use the same sample of 618,945 loans and include county and property-type fixed effects, loan characteristics, initial loan-stock controls, and WFH exposure. Column 2 adds parent-bank fixed effects, and column 3 adds local controls and the peer resolution rate. Resolution counts and WFH exposure are standardized by its standard deviation separately, same as the previous two prediction tests.

The results show that the proportional response to earlier resolutions is consistently larger in more WFH-exposed markets. Across all specifications, the interaction between earlier resolutions and WFH exposure is positive and statistically significant. The estimate is 0.124 in the baseline, 0.123 after adding parent-bank fixed effects, and 0.128 in the full specification. The interpretation of the magnitude of these coefficients differs from that in the previous tables. In the full specification, at zero standardized WFH exposure, a one-standard-deviation increase in earlier resolution counts is associated with a $100[\exp(0.284) - 1] \approx 32.8$ percent increase in subsequent resolution probability. At WFH exposure one standard deviation higher, the corresponding increase is $100[\exp(0.284 + 0.128) - 1] \approx 51$ percent. The estimates are consistent with the model's prediction of a larger proportional response to information from earlier resolutions where initial valuation uncertainty is greater.

As a separate test, I also estimate an OLS specification to align with the first two tests. More precisely, I regress the extension indicator on earlier resolution counts and their interaction with WFH exposure, using the same specifications as before. This regression provides additional evidence on extension responses in levels. With full controls, a one-standard-deviation increase in WFH exposure makes the extension response to approximately 113 earlier resolutions 0.099 percentage points more negative ($p = 0.064$). Earlier resolutions are thus associated with a larger decline in

Table 5: Initial Uncertainty and Proportional Resolution Responses

	(1) Baseline	(2) + Bank FE	(3) + Local controls and rate
Resolution counts × WFH	0.124***	0.123***	0.128***
	(0.030)	(0.030)	(0.030)
Resolution counts (SD)	0.202***	0.220***	0.284***
	(0.057)	(0.056)	(0.062)
WFH (SD)	−0.016	−0.014	−0.013
	(0.070)	(0.063)	(0.062)
Resolution rate R (pp)			−0.005
			(0.003)
Parent-bank fixed effects	No	Yes	Yes
Local controls	No	No	Yes
Initial peer-stock controls	Yes	Yes	Yes
Peer resolution rate	No	No	Yes
Observations	618,945	618,945	618,945

Notes. Outcome: default or sale payoff during the later window. All regressions are estimated using Poisson pseudo-maximum likelihood model. Earlier peer events occur during 2020Q1–2021Q4 and exclude the loan's own parent bank. All count regressors are standardized by the common total-event SD of 112.9793; rates are in unit of percentage points. All columns include county and property-type fixed effects, log loan age, origination LTV, raw minimum quarters to maturity, and standardized WFH exposure. Local controls are pre-pandemic turnover, log transaction volume, and bank concentration. Counts are uncentered. All columns include initial peer-stock controls; column 3 adds the peer rate. Standard errors are clustered at county level, and are noted in parenthness. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

extension in more WFH-exposed markets. Appendix D.3 reports the detailed results.

5.4 Ruling out Alternative Explanations: Direction of Price News

I next use the direction of transaction-price news to distinguish uncertainty reduction from an explanation based on loss recognition.

A possible alternative explanation draws on the logic of Gorton and Ordoñez (2014). Their framework shows how a small shock can develop into a large financial crisis by inducing lenders to acquire information about collateral quality. In normal times, credit can flow without detailed information about the underlying collateral. A negative macroeconomic shock, however, can prompt lenders to pay the cost of examining collateral quality. The resulting information production can disrupt lending and trigger a "collateral crisis."

In the present setting, a related channel could operate through regulatory scrutiny. Suppose banks continue to extend distressed loans partly because regulators cannot precisely assess the underlying collateral. Low transaction prices could challenge the valuations used to justify those extensions. Regulators may then require more conservative valuations, greater loss recognition, or corrective action, making further extensions less attractive. Earlier resolutions would therefore be

followed by fewer extensions, even if the information they reveal does not reduce banks' uncertainty about the gains from resolving their loans. Under this explanation, better information about collateral allows regulators to constrain banks' continued extension of distressed loans.

With additional assumptions about how regulators update their beliefs, this explanation could also lead to similar heterogeneity. Public price disclosure could strengthen the extension response by making evidence of lower collateral values more accessible to regulators. Greater initial uncertainty could also make new transaction evidence more important for regulatory assessments. Explaining Prediction 3 under this alternative explanation is more demanding, however, because it requires additional assumptions about how regulators update collateral values and why these updates would lead to a larger proportional increase in resolution when initial uncertainty is greater.

One way to distinguish the proposed mechanism from this alternative explanation is to examine whether transactions convey good or bad price news. Under the "regulator learning" explanation, transactions conveying bad price news reveal collateral shortfalls and make continued extension harder to justify. Transactions conveying good price news do not create the same pressure to recognize losses; instead, they may support existing collateral valuations, making it easier for banks to extend credit on risky assets. Thus, if this alternative channel explains the negative relationship between earlier sales and subsequent extension, that relationship should be concentrated among transactions conveying bad price news. Yet the public-learning mechanism is different. In our model, the resolution cutoff depends on public belief precision rather than the posterior mean of collateral value. Both good- and bad-news resolution outcomes reduce valuation uncertainty and the incentive to wait for further evidence. I therefore examine whether the negative extension response shows up with both types of news from resolution events. More precisely, I test whether the negative relationship between extension and prior resolution remains when restricting attention to resolutions conveying good news.

5.4.1 Measuring Price News and Empirical Design

I use transaction prices to distinguish good from bad price news revealed by sale payoffs. The test uses sale-payoff events in disclosure states and requires an observed transaction price and a prior assessed value. For each transaction, I compare the sale price with an assessed-value benchmark adjusted to the transaction date. The assessed value is one of the key recorded values that regulators and appraisers can examine when assessing collateral values, thus it provides a natural benchmark for me to measure the sign of the transaction price news.

However, one concern is that assessed values can be staled. Tax authorities may be slow to update property values to reflect market changes and may continue to use values based on earlier transactions. As a result, the gap between the transaction price and the assessed value may be positive even when the market interprets the transaction price as bad news. The gap may also

depend on when the assessed value was last updated. For example, a property whose assessed value reflects a transaction in 2008 may not be directly comparable to one whose assessed value reflects a transaction in 2018, making the raw gaps less informative about whether the news is good or bad.

I address this in two steps. First, I identify the most recent transaction at which the assessed value was aligned with the transaction price, then roll that value forward using regional and property-type price indices. Second, I classify news using the price gap relative to its pooled mean. Events above the mean are good news (G), and those at or below it are bad news (B). This cutoff accounts for the average difference between transaction prices and the adjusted assessed values. The two groups distinguish relatively high from low transaction valuations, allowing me to test whether the extension response is stronger after bad price news.

I aggregate the classified sale payoffs within each county–property-type market during the earlier period, excluding the loan's own parent bank. I then replace the total sale-payoff count with separate good- and bad-news counts:

$$\begin{aligned} Ext_{i,c,p,b,t_2} = \alpha + \beta_G SaleCounts^{G}_{c,p,-b,t_1} + \beta_B SaleCounts^{B}_{c,p,-b,t_1} \\ + \Theta' W_{i,c,p,b,t_0} + \lambda_c + \kappa_p + \psi_b + \varepsilon_{i,c,p,b,t_2}. \end{aligned}$$

Both counts enter jointly and are divided by the same standard deviation of total sale-payoff counts. Thus, the two coefficients are interpreted by describe the extension response to an equal increase in the number of transactions. The dependent variable, controls, and fixed effects are the same as in the previous tests. Standard errors are clustered by county.

The coefficients β_G and β_B describe the relationships between subsequent extension and earlier sale payoffs conveying relatively good and relatively bad price news, respectively. The price-discovery mechanism predicts that both coefficients are negative because, regardless of the realized value, resolution events that provide information about collateral values can reduce uncertainty and, therefore, lead to less waiting.

5.4.2 Results

Table 6 reports the results for 544,506 loans in disclosure states. All columns use the same sample and include loan and local controls, initial loan-stock controls, and county, property-type, and parent-bank fixed effects. Column 1 relates subsequent extensions to all earlier sale payoffs at other banks. Column 2 adds standardized mean transaction-price news, while column 3 separates price-observed sale payoffs into good- and bad-news counts. All count variables use the same sale-count standard deviation, allowing their coefficients to be compared on a common scale. Appendix D.4 provides the event-coverage and scaling details.

Table 6: Peer Sale Payoffs Conveying Good and Bad Price News

	(1) All sales	(2) + Mean news	(3) Good / bad
All peer sale payoffs	-0.668^{***}	-0.656^{***}	
	(0.140)	(0.138)	
Mean transaction price news		-0.054	
		(0.062)	
Good-news sale payoffs			-0.671^{***}
			(0.161)
Bad-news sale payoffs			-0.691^{**}
			(0.276)
Loan and local controls	Yes	Yes	Yes
County, property, and bank FE	Yes	Yes	Yes
Initial peer-stock controls	Yes	Yes	Yes
Disclosure states only	Yes	Yes	Yes
Observations	544,506	544,506	544,506
R^2	0.0722	0.0722	0.0722
Good = bad: p-value			0.939

Notes. Outcome: extension during 2022Q1–2025Q2. Coefficients and standard errors are in percentage points. All counts use the common all-sale-count SD of 78.4594. Good news is above the pooled mean of the transaction-price gap; bad news is at or below it. All columns include county and property-type fixed effects, log loan age, origination LTV, raw minimum quarters to maturity, and standardized WFH exposure. Local controls are pre-pandemic turnover, log transaction volume, and bank concentration. All columns include initial peer-stock controls. Standard errors clustered by county are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Controlling for average price news leaves the sale-count coefficient nearly unchanged. In column 1, we see the coefficient is −0.668. After controlling for the average price news, we see the coefficient remains similar at −0.656, both significant at the 1 percent level. In addition, the coefficient on mean price news is small and statistically insignificant at −0.054. That is, the negative sale-count relationship persists after accounting for whether transactions convey relatively high or low valuations on average, suggesting that it is not explained solely by the direction of measured price news.

Column 3 shows that both good- and bad-news transactions are followed by less extension. The coefficient on relatively good price news is −0.671, significant at the 1 percent level, while the bad-news coefficient is −0.691, significant at the 5 percent level. An increase in either category equal to one standard deviation of total sale-payoff counts is therefore associated with an approximately 0.67–0.69 percentage-point reduction in subsequent extension probability. The two point estimates are similar, indicating that an equal increase in the number of good- or bad-news sale payoffs is associated with a similar magnitude of decline in extension probability.

These results help distinguish the proposed mechanism from the alternative explanation. If earlier sales discouraged extension exclusively by revealing collateral losses to the public and regulators, the negative relationship should be concentrated among transactions conveying bad

price news, or we should see a stronger response to relatively bad news in the regression. In the regression results, we see that transactions conveying relatively good price news are also associated with less extension, with a more precisely estimated coefficient compared with the bad price new events. Additionally, the sale-count coefficient remains stable after controlling for mean price news. These findings support the interpretation that transactions matter beyond revealing low collateral values. Both good- and bad-news outcomes can reduce valuation uncertainty and the incentive to wait, as emphasized by the model. The evidence thus challenges the alternative explanation based exclusively on adverse price news.

6 Conclusion

This paper explains how limited trading and delayed loan resolution can coexist after a lasting sectoral shock. In the model, a bank that resolves a distressed loan sells its collateral, where the price includes the uncertainty discount. The key assumption is that each resolution generates public information to inform the collateral value, which benefits lenders whose loans remain outstanding. Each bank therefore has an incentive to keep its loan while waiting to learn from others' resolutions. When many banks wait, few transactions occur, slowing the production of information about collateral values. Completed resolutions reduce uncertainty and make subsequent resolutions more attractive, creating complementarity across periods.

The loan-level evidence from the post-pandemic CRE market aligns well with this mechanism. I start by showing that earlier resolutions at other banks are associated with lower subsequent extension probabilities. This relationship persists under 2SLS, using a shift-share instrument that relies on the exogeneity of local exposure to bank-level shocks. Next, I show that the reduction in extension probabilities following sale payoffs is larger where transaction prices are publicly disclosed. Finally, I show that earlier resolutions are associated with a larger proportional increase in subsequent resolution probability in markets that are more exposed to the WFH shock. These three tests follow closely from the model's predictions and support the mechanism through which waiting for others to resolve can allow uncertainty and delay to persist.

To address concerns about alternative mechanisms, I additionally show that both relatively good and bad price news are followed by fewer extensions, suggesting that the relationship is not confined to transactions that reveal lower collateral values. These findings support the view that information from completed resolutions reduces uncertainty and encourages other lenders to resolve their outstanding loans.

The paper offers an explanation for delayed resolution that complements the conventional zombie-lending story about risk-shifting or information asymmetry . In the model, information free-riding can sustain loan extensions even at well-capitalized banks. The empirical analysis provides novel tests of social learning in the CRE market after the WFH shock by examining how

lenders respond to the availability and content of transaction information. The mechanism can apply beyond CRE, in lending markets where sectoral shocks leave collateral values uncertain, and sales provide useful information to other lenders.

The main message of this paper is that information about collateral values can shape how credit markets adjust to lasting sectoral shocks. Extending a distressed loan prolongs the lender's exposure and, at the same time, can postpone a transaction that would help other lenders value their collateral. When information is a public good, this delay can lead to prolonged endogenous waiting. Policies that affect transaction activity or public access to transaction information may therefore influence resolution beyond the loans they directly affect. Future research could evaluate these spillovers, weighing the information benefits of early resolutions against the costs of accelerating sales. This could clarify when facilitating price discovery improves the allocation of assets and credit during potential crises and helps prevent financial instability.

References

Adão, R., Kolesár, M., & Morales, E. (2019). Shift-share designs: Theory and inference. *The Quarterly Journal of Economics*, *134*(4), 1949–2010. https://doi.org/10.1093/qje/qjz025

Adão, R., Kolesár, M., Morales, E., & Zhang, X. (2020). *IVREG_SS: Stata module to compute confidence intervals, standard errors, and p-values in an IV regression in which the excluded instrumental variable has a shift-share structure.* Statistical Software Components S458759, Boston College Department of Economics. https://ideas.repec.org/c/boc/bocode/s458759.html

Anenberg, E., Kim, Y. S., & Moszkowski, E. (2025). *Work-from-home, commercial real estate risk and credit supply: Evidence from a large sample of bank loan portfolios.* https://www.aeaweb.org/conference/2026/program/paper/y39fz3iS (Working paper, September 5)

BankUnited, Inc. (2016). *2016 annual report* (Tech. Rep.). BankUnited, Inc. https://www.sec.gov/Archives/edgar/data/1504008/000150400817000004/document-20161231.htm (Form 10-K for the year ended December 31, 2016)

Barber, B. M., & Yasuda, A. (2017). Interim fund performance and fundraising in private equity. *Journal of Financial Economics*, *124*(1), 172–194. https://www.ayakoyasuda.com/papers/2017/Paper12JFE.pdf

Barrero, J. M., Bloom, N., & Davis, S. J. (2023). The evolution of work from home. *Journal of Economic Perspectives*, *37*(4), 23–49. https://doi.org/10.1257/jep.37.4.23

Berrens, R. P., & McKee, M. (2004). What price nondisclosure? the effects of nondisclosure of real estate sales prices. *Social Science Quarterly*, *85*(2), 509–520. https://onlinelibrary.wiley.com/doi/abs/10.1111/j.0038-4941.2004.08502017.x

Birch, R. (2018). Where are NYC taxi medallion values heading? *CU Today*. https://www.cutoday.info/THE-feature/Where-Are-NYC-Taxi-Medallion-Values-Heading (January 22)

Bruche, M., & Llobet, G. (2014). Preventing zombie lending. *The Review of Financial Studies*, *27*(3), 923–956. https://doi.org/10.1093/rfs/hht064

Caballero, R. J., Hoshi, T., & Kashyap, A. K. (2008). Zombie lending and depressed restructuring in Japan. *American Economic Review*, *98*(5), 1943–1977. https://doi.org/10.1257/aer.98.5.1943

Callaway, B., & Sant'Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of econometrics*, *225*(2), 200–230.

Cao, J., Song, L., Titman, S., & Zhan, X. (2025). *Real activities and uncertainty: Evidence from real estate markets* (Working Paper No. 04/2025). Hong Kong Institute for Monetary and Financial Research. https://www.aof.org.hk/docs/default-source/hkimr/working-papers/2025/full-text_wp-no-04_2025.pdf

CBRE. (2024). *U.S. cap rate survey H1 2024* (Tech. Rep.). CBRE. https://www.cbre.com/insights/reports/14195699984 (July 31)

Chamley, C., & Gale, D. (1994). Information revelation and strategic delay in a model of investment. *Econometrica*, *62*(5), 1065–1085. https://doi.org/10.2307/2951507

Crosignani, M., & Prazad, S. (2024). *Extend-and-pretend in the U.S. CRE market* (Staff Reports

No. 1130). Federal Reserve Bank of New York. https://www.newyorkfed.org/research/staff_reports/sr1130 (Revised June 2026)

Dingel, J. I., & Neiman, B. (2020). How many jobs can be done at home? *Journal of Public Economics*, *189*, 104235. https://doi.org/10.1016/j.jpubeco.2020.104235

Drechsler, I., Savov, A., & Schnabl, P. (2021). Banking on deposits: Maturity transformation without interest rate risk. *The Journal of Finance*, *76*(3), 1091–1143. https://onlinelibrary.wiley.com/doi/10.1111/jofi.13013

Ellis, L., & Naughtin, C. (2010). Commercial property and financial stability: An international perspective. *Reserve Bank of Australia Bulletin*. https://www.rba.gov.au/publications/bulletin/2010/jun/4.html (June)

Fajgelbaum, P. D., Schaal, E., & Taschereau-Dumouchel, M. (2017). Uncertainty traps. *The Quarterly Journal of Economics*, *132*(4), 1641–1692. https://doi.org/10.1093/qje/qjx021

Federal Financial Institutions Examination Council. (2025). Instructions for preparation of consolidated reports of condition and income: FFIEC 031 and FFIEC 041 [Computer software manual]. https://www.ffiec.gov/sites/default/files/data/reporting-forms/FFIEC031_FFIEC041_202503_i.pdf (Updated March 2025)

Federal Reserve Board. (2023). *Policy statement on prudent commercial real estate loan accommodations and workouts.* Joint statement with the Federal Deposit Insurance Corporation, National Credit Union Administration, and Office of the Comptroller of the Currency. https://www.federalreserve.gov/supervisionreg/srletters/SR2305.htm (June 29)

Gárate, S., Hilterbrand, C., & Pennington-Cross, A. (2025). *The impact of non-disclosure laws on the single-family housing market: Appraisal bias and mortgage default*. https://realestatecenter.business.fsu.edu/sites/g/files/upcbnu2141/files/Non_disclosure_paper_feb_2025.pdf (Working paper, February 2025 version)

Glancy, D. (2026). *Pretend or amend? on evergreening in CRE* (Finance and Economics Discussion Series No. 2026-025). Board of Governors of the Federal Reserve System. https://www.federalreserve.gov/econres/feds/files/2026025pap.pdf

Glancy, D., Kurtzman, R. J., & Loewenstein, L. (2022). *Loan modifications and the commercial real estate market* (Finance and Economics Discussion Series No. 2022-050). Board of Governors of the Federal Reserve System. https://www.federalreserve.gov/econres/feds/files/2022050pap.pdf

Goldsmith-Pinkham, P., Sorkin, I., & Swift, H. (2020). Bartik instruments: What, when, why, and how. *American Economic Review*, *110*(8), 2586–2624. https://www.aeaweb.org/articles?id=10.1257/aer.20181047

Gorton, G., & Ordoñez, G. (2014). Collateral crises. *American Economic Review*, *104*(2), 343–378. https://www.aeaweb.org/articles?id=10.1257/aer.104.2.343

Greenstone, M., Mas, A., & Nguyen, H.-L. (2020). Do credit market shocks affect the real economy? quasi-experimental evidence from the Great Recession and "normal" economic times. *American Economic Journal: Economic Policy*, *12*(1), 200–225. https://doi.org/10.1257/pol.20160005

Gupta, A., Mittal, V., & Van Nieuwerburgh, S. (2026). Work from home and the office real estate apocalypse. *American Economic Review*, *116*(2), 674–709. https://doi.org/10.1257/aer.20231619

Harsanyi, J. C. (1973). Games with randomly disturbed payoffs: A new rationale for mixed-strategy equilibrium points. *International Journal of Game Theory*, *2*, 1–23. https://link.springer.com/article/10.1007/BF01737554

Hinzen, F. J., Severino, F., & Van Nieuwerburgh, S. (2026). *Too-many-to-ignore: Regional banks and CRE risks.* https://franzhinzen.com/assets/pdf/HSV%20-%20Regional%20Banks%20and%20CRE%20Risks.pdf (Working paper, January)

Hu, Y., & Varas, F. (2021). A theory of zombie lending. *The Journal of Finance*, *76*(4), 1813–1867. https://doi.org/10.1111/jofi.13022

Hudgins, C. (2024). US office market report: Higher rates, lower valuations spur fierce headwinds. *S&P Global Market Intelligence*. https://www.spglobal.com/market-intelligence/en/news-insights/research/us-office-market-report-higher-rates-lower-valuations-spur-fierce-headwinds (January 22)

Jiang, E. X., Matvos, G., Piskorski, T., & Seru, A. (2025). Monetary tightening, commercial real estate distress, and US bank fragility. *Journal of Political Economy Macroeconomics*, *3*(4), 525–573. https://doi.org/10.1086/738234

Kurlat, P. (2018). Liquidity as social expertise. *The Journal of Finance*, *73*(2), 619–656. https://doi.org/10.1111/jofi.12606

Ling, D. C., Naranjo, A., & Scheick, B. (2021). There is no place like home: Information asymmetries, local asset concentration, and portfolio returns. *Real Estate Economics*, *49*(1), 36–74. https://onlinelibrary.wiley.com/doi/full/10.1111/1540-6229.12336

Office of the Legislative Auditor General. (2024, April). *A performance audit of Utah's property tax system: A case for increased accountability and transparency* (Report No. 2024-05). Utah Legislature. https://le.utah.gov/interim/2024/pdf/00001999.pdf

Peters, A. (2017). The slow remaking of the taxi-medallion loan market. *American Banker*. https://www.americanbanker.com/news/the-slow-remaking-of-the-taxi-medallion-loan-market (May 25)

Rajan, R. G., & Zingales, L. (1998). Financial dependence and growth. *American Economic Review*, *88*(3), 559–586. https://www.nber.org/papers/w5758

Veldkamp, L. L. (2005). Slow boom, sudden crash. *Journal of Economic Theory*, *124*(2), 230–257. https://doi.org/10.1016/j.jet.2003.12.008

Appendix for

When Banks Wait: Strategic Delay for Distressed Loan Resolution

Judy Yue
University of Chicago

November 13, 2026

A Model Proofs and Additional Results **IA.2**
A.1 Equivalent Lagged-State Recursive Representation . . . IA.2
A.2 A Fixed-Waiting-Cost Benchmark . . . IA.3
A.3 Measuring and Computing Resolution Delays . . . IA.5
A.4 Separation of the Posterior Mean . . . IA.7
A.5 Option Value and Equilibrium Construction . . . IA.8
A.6 Proofs of the Strategic-Waiting Results . . . IA.11
A.7 Proof of Across-Period Complementarity . . . IA.12
A.8 Baseline Comparative Statics . . . IA.15
A.9 Proofs of the Amplification Results . . . IA.22
A.10 Proofs for the Fixed-Waiting-Cost Benchmark . . . IA.28
A.11 Additional Numerical Illustrations . . . IA.32

B Additional Data Documentation **IA.37**
B.1 Lender and Property Linkage . . . IA.37
B.2 Loan Identity, Boundaries, and Quarterly Histories . . . IA.38
B.3 Event Classification and Outcome Aggregation . . . IA.38
B.4 Cohort, Peer Exposure, and Estimation Restrictions . . . IA.39
B.5 WFH Exposure and Control Variables . . . IA.40
B.6 Disclosure and Relative Transaction-Price News . . . IA.41
B.7 Coverage Comparison . . . IA.41
B.8 CompStak Coverage . . . IA.42

C Additional Background and Motivating Evidence **IA.43**
C.1 Office Leasing Revenue . . . IA.43
C.2 Descriptive Extension Shares . . . IA.43
C.3 Alternative Capital-Buffer Split . . . IA.44
C.4 GSIBs and Banks in the Stress-Tested Sample . . . IA.45

D Additional Empirical Results **IA.47**
D.1 Sensitivity to the Window Cutoff . . . IA.47
D.2 Joint and Default-Only Disclosure Comparisons . . . IA.47
D.3 Initial Uncertainty and Extension in Levels . . . IA.50
D.4 Price-News Coverage and Scaling . . . IA.51

A Model Proofs and Additional Results

A.1 Equivalent Lagged-State Recursive Representation

The lagged-state representation records the transition that produced the current public state. Let

$$\widehat{S}_{t+1} = (\mu_t, \gamma_t, N_t, n_t, X_t),$$

where (μ_t, γ_t, N_t) is the public state at the beginning of period t, n_t is the number of resolutions that occurred, and X_t is their average signal, with X_t a placeholder when $n_t = 0$. For an active current state, $n_t \in \{0, \dots, N_t - 1\}$. The current public state $S = S_{t+1} = (\mu, \gamma, N)$ is recovered according to

$$\gamma = \gamma_t + n_t\delta, \qquad N = N_t - n_t, \qquad \mu = \frac{\gamma_t\mu_t + n_t\delta X_t}{\gamma_t + n_t\delta}, \tag{A.1}$$

with $\mu = \mu_t$ when $n_t = 0$. Thus, the lagged representation is an augmented description of the transition; all equilibrium objects factor through the current state S.

Definition. RCE in lagged variables. An equivalent symmetric recursive equilibrium consists of a strategy s^*, its associated resolution probability $q^* = q^{s^*}$, value functions V and $\mathcal{V}$, the price function P, and a law of motion such that, for every lagged representation $\widehat{S}_{t+1}$ and current cost c, with the current state $S = (\mu, \gamma, N)$ determined by (A.1):

1. (*Bellman system*) The value functions satisfy

$$V(S, c) = \max\left\{P(S), -c + \sum_{n=0}^{N-1} \binom{N-1}{n} q^*(S)^n (1 - q^*(S))^{N-1-n} \mathbb{E}[\mathcal{V}(S'(n)) \mid S, n]\right\}, \tag{A.2}$$

$$\mathcal{V}(S) = \int V(S, c)\, dG(c), \tag{A.3}$$

where, when $n \geq 1$, the next current state is

$$S'(n) = \left(\frac{\gamma\mu + n\delta X'}{\gamma + n\delta}, \gamma + n\delta, N - n\right),$$

and $S'(0) = S$. The expectation is over X'.

2. (*Cutoff and consistency*) The equilibrium strategy satisfies

$$s^*(S, c) = \mathbf{1}\{c \geq c^*(S)\},$$

where

$$c^*(S) = \frac{a}{2\gamma} + \frac{a}{2\delta} + \sum_{n=0}^{N-1} \binom{N-1}{n} q^*(S)^n (1 - q^*(S))^{N-1-n} \Big(\mathbb{E}[\mathcal{V}(S'(n)) \mid S, n] - \mu\Big), \tag{A.4}$$

and

$$q^*(S) = 1 - G(c^*(S)),$$

with G extended below and above its support in the usual way.

3. (*Price*) The resolution price is

$$P(S) = \mu - \frac{a}{2\gamma} - \frac{a}{2\delta}.$$

4. (*Aggregate law of motion*) Given the current state S, the number of current resolutions satisfies

$$n' \sim \text{Binomial}(N, q^*(S)).$$

Conditional on $n' \geq 1$, their average signal has predictive distribution

$$X' \mid (S, n') \sim N\left(\mu, \frac{1}{\gamma} + \frac{1}{n'\delta}\right). \tag{A.5}$$

If at least one bank remains, next period's lagged representation is (μ, γ, N, n', X'). In bank i's Bellman equation, the same kernel is applied to its $N-1$ peers because bank i remains unresolved after choosing to wait.

This representation is equivalent to the symmetric Markov equilibrium in the main text because $\widehat{S}_{t+1}$ affects prices, payoffs, and strategies only through the sufficient current state $S = (\mu, \gamma, N)$.

A.2 A Fixed-Waiting-Cost Benchmark

To illustrate the strategic mechanism, consider a benchmark in which every bank faces the same deterministic waiting cost in each period:

$$c_{i,t} \equiv \kappa > 0.$$

This replaces the baseline's continuous cost distribution with a common, known cost. Without idiosyncratic cost heterogeneity, an interior symmetric equilibrium generally requires banks to randomize between resolution and waiting. Let

$$q^\kappa(\gamma, N; \delta, a)$$

denote the probability with which each unresolved bank resolves, where randomization is independent across banks.

Write the total resolution discount as

$$d(\gamma) \equiv \frac{a}{2\gamma} + \frac{a}{2\delta}.$$

If each of the other $N-1$ banks resolves with probability q, let

$$\begin{aligned} \mathcal{I}_{\gamma,N}(q; \delta, a) &\equiv \mathbb{E}\left[d(\gamma) - d(\gamma + J\delta)\right], \qquad J \sim \text{Binomial}(N-1, q) \\ &= \sum_{j=0}^{N-1} \binom{N-1}{j} q^j (1-q)^{N-1-j} \left[d(\gamma) - d(\gamma + j\delta)\right]. \end{aligned} \tag{A.6}$$

This is the expected reduction in the uncertainty discount obtained by waiting for information

generated by peer resolutions.

Proposition 7. Equilibrium with a fixed waiting cost.

At every active state, there exists a unique symmetric independently mixed Markov equilibrium resolution probability. When $N = 1$,

$$q^{\kappa}(\gamma, 1; \delta, a) = 1.$$

When $N \geq 2$,

$$q^{\kappa}(\gamma, N; \delta, a) = \begin{cases} 1, & \kappa \geq d(\gamma) - d(\gamma + (N-1)\delta), \\ \text{the unique } q \in (0,1) \text{ satisfying } \mathcal{I}_{\gamma,N}(q; \delta, a) = \kappa, & \kappa < d(\gamma) - d\big(\gamma + (N-1)\delta\big). \end{cases} \tag{A.7}$$

The equilibrium value of an unresolved bank is

$$v^{\kappa}(\gamma, N) = -d(\gamma). \tag{A.8}$$

Proof: See Appendix A.10.1 for the full proof.

The scalar characterization makes the two strategic forces especially clear.

Proposition 8. Strategic substitution and across-period complementarity.

In the fixed-cost benchmark:

1. *Within-period strategic substitution.* A higher resolution probability among the other banks strictly increases the payoff advantage of waiting:

$$\frac{\partial}{\partial q}\left[\mathcal{I}_{\gamma,N}(q; \delta, a) - \kappa\right] > 0. \tag{A.9}$$

   Thus, a higher peer-resolution probability makes the bank weakly less willing to resolve.

2. *Across-period complementarity.* For every $N \geq 2$,

$$q^{\kappa}(\gamma + \delta, N-1; \delta, a) \geq q^{\kappa}(\gamma, N; \delta, a). \tag{A.10}$$

   The inequality is strict whenever

$$q^{\kappa}(\gamma, N; \delta, a) < 1.$$

   Thus, one realized peer resolution weakly raises the subsequent resolution probability of each remaining bank.

Proof: See Appendix A.10.2 for the full proof.

The fixed-cost characterization also simplifies the amplification results. Define the effect of one realized peer resolution as

$$\Delta q^{\kappa}(\gamma, N; \delta, a) \equiv q^{\kappa}(\gamma + \delta, N-1; \delta, a) - q^{\kappa}(\gamma, N; \delta, a). \tag{A.11}$$

Proposition 9. Amplification in the fixed-cost benchmark.

1. *Signal informativeness.* Fix (γ, N, a, κ), with $N \geq 2$, and suppose

$$\delta^H > \delta^L > 0, \qquad \delta^H \delta^L \geq \frac{\gamma^2}{2}.$$

Then

$$\Delta q^\kappa(\gamma, N; \delta^H, a) \geq \Delta q^\kappa(\gamma, N; \delta^L, a). \tag{A.12}$$

2. *Current public precision.* Fix $N \geq 3$, $\delta > 0$, and $0 < \gamma^L < \gamma^H$. For sufficiently large a,

$$\Delta \log q^\kappa(\gamma^L, N; \delta, a) > \Delta \log q^\kappa(\gamma^H, N; \delta, a), \tag{A.13}$$

where

$$\begin{aligned}\Delta \log q^\kappa(\gamma, N; \delta, a) \equiv\ & \log q^\kappa(\gamma + \delta, N - 1; \delta, a) \\ & - \log q^\kappa(\gamma, N; \delta, a).\end{aligned}$$

Moreover,

$$\Delta \log q^\kappa(\gamma, N; \delta, a) \longrightarrow \log\left(\frac{N-1}{N-2}\right) + \log\left(1 + \frac{2\delta}{\gamma}\right) \tag{A.14}$$

as $a \to \infty$.

Proof: See Appendix A.10.3 for the full proof.

Remark. When $N = 2$, the equilibrium probability has the closed form

$$q^\kappa(\gamma, 2; \delta, a) = \min\left\{1, \frac{\kappa}{d(\gamma) - d(\gamma + \delta)}\right\}. \tag{A.15}$$

After one peer resolution, only one bank remains and resolves with probability one. Hence

$$\Delta q^\kappa(\gamma, 2; \delta, a) = 1 - q^\kappa(\gamma, 2; \delta, a).$$

In the two-bank case, a more valuable peer signal lowers the initial resolution probability, while a completed peer resolution removes the remaining incentive to wait.

A.3 Measuring and Computing Resolution Delays

Fix an active initial reduced state (γ_0, N_0), with $N_0 \geq 1$. For each bank $i \in \{1, \ldots, N_0\}$ that is unresolved at the beginning of period zero, recall that $d_{i,t} = s^*(S_t, c_{i,t}) = 1$ when it resolves. Define its resolution time by

$$\tau_i \equiv \inf\{t \geq 0 : d_{i,t} = 1\}.$$

Thus, $\tau_i = 0$ if bank i resolves immediately. Define the time until all initially unresolved banks have resolved by

$$\tau^{\text{all}} \equiv \max_{1 \leq i \leq N_0} \tau_i = \inf\{t \geq 0 : N_{t+1} = 0\}.$$

By symmetry, all initially unresolved banks have the same expected resolution time. Using the tail-sum formula and

$$N_t = \sum_{i=1}^{N_0} \mathbf{1}\{\tau_i \geq t\},$$

expected individual resolution delay can be written as

$$\begin{aligned} D^{\text{ind}}(\gamma_0, N_0; \delta) &\equiv \frac{1}{N_0} \sum_{i=1}^{N_0} \mathbb{E}[\tau_i \mid \gamma_0, N_0] \\ &= \frac{1}{N_0} \sum_{i=1}^{N_0} \sum_{t=1}^{\infty} \Pr(\tau_i \geq t \mid \gamma_0, N_0) \\ &= \sum_{t=1}^{\infty} \mathbb{E}\left[\frac{N_t}{N_0} \middle| \gamma_0, N_0 \right]. \end{aligned} \tag{A.16}$$

Similarly, the expected duration of the entire resolution process is

$$\begin{aligned} D^{\text{all}}(\gamma_0, N_0; \delta) &\equiv \mathbb{E}[\tau^{\text{all}} \mid \gamma_0, N_0] \\ &= \sum_{t=1}^{\infty} \Pr(\tau^{\text{all}} \geq t \mid \gamma_0, N_0) \\ &= \sum_{t=1}^{\infty} \Pr(N_t > 0 \mid \gamma_0, N_0). \end{aligned} \tag{A.17}$$

Thus, D^{ind} is the discrete area under the expected fraction of initially unresolved banks remaining in Figure 4, while D^{all} measures the expected duration of the entire resolution process. Both are finite: the set of states reachable from (γ_0, N_0) is finite, and every active state has a strictly positive probability of at least one resolution.

Fix the signal technology and the other primitives. At an active state (γ, N), the probability of n resolutions is

$$p_n(\gamma, N) = \binom{N}{n} q^*(\gamma, N)^n [1 - q^*(\gamma, N)]^{N-n}, \qquad n = 0, \ldots, N.$$

Let $T(\gamma, N)$ be the expected number of periods with an active market, including the current period. Let $R(\gamma, N)$ be the expected sum of unresolved-bank counts over those periods. Set $T(\gamma, 0) = R(\gamma, 0) = 0$. Separating the no-resolution transition, which leaves the state unchanged, gives the backward recursions

$$\begin{aligned} T(\gamma, N) &= \frac{1 + \sum_{n=1}^{N} p_n(\gamma, N) T(\gamma + n\delta, N - n)}{1 - p_0(\gamma, N)}, \\ R(\gamma, N) &= \frac{N + \sum_{n=1}^{N} p_n(\gamma, N) R(\gamma + n\delta, N - n)}{1 - p_0(\gamma, N)}. \end{aligned}$$

The denominator is positive because $q^*(\gamma, N) > 0$. With the convention above that immediate

resolution has zero delay,

$$D^{\text{all}}(\gamma, N) = T(\gamma, N) - 1, \qquad D^{\text{ind}}(\gamma, N) = \frac{R(\gamma, N)}{N} - 1.$$

These recursions compute the survival-sum measures in (A.16) and (A.17) used in the numerical delay figures.

A.4 Separation of the Posterior Mean

Proof of Lemma 1. Conjecture that

$$\mathcal{V}(S) = \mu + v(\gamma, N).$$

Conditional on any number n of peer resolutions, the posterior mean remains a martingale:

$$\mathbb{E}[\mu' \mid S, n] = \mu.$$

This follows because resolution decisions are driven by waiting-cost draws that are independent of θ and of the resolution signals.

Substituting the conjectured form into the definition of $V(S, c)$ gives

$$\begin{aligned} V(S, c) &= \max\left\{\mu - \frac{a}{2\gamma} - \frac{a}{2\delta}, -c + \mathbb{E}_{n,X'}\left[\mu' + v(\gamma', N')\right]\right\} \\ &= \mu + \max\left\{-\frac{a}{2\gamma} - \frac{a}{2\delta}, -c + \mathbb{E}_n\left[v(\gamma', N')\right]\right\}. \end{aligned}$$

Taking expectations over the current waiting-cost draw therefore yields

$$\mathcal{V}(S) = \mu + \mathbb{E}_c\left[\max\left\{-\frac{a}{2\gamma} - \frac{a}{2\delta}, -c + \mathbb{E}_n\left[v(\gamma', N')\right]\right\}\right].$$

Hence the reduced value satisfies

$$v(\gamma, N) = \mathbb{E}_c\left[\max\left\{-\frac{a}{2\gamma} - \frac{a}{2\delta}, -c + \mathbb{E}_n\left[v(\gamma + n\delta, N - n)\right]\right\}\right].$$

At the cutoff, the bank is indifferent between resolving and waiting:

$$-\frac{a}{2\gamma} - \frac{a}{2\delta} = -c^*(S) + \mathbb{E}_n\left[v(\gamma', N')\right].$$

Therefore,

$$c^*(\gamma, N) = \frac{a}{2\gamma} + \frac{a}{2\delta} + \mathbb{E}_n\left[v(\gamma', N')\right].$$

The Bellman and consistency conditions therefore preserve the affine form, and the cutoff and resolution probability depend only on (γ, N). The reduced system in the main text constructs the equilibrium; Theorem 2 establishes its existence and uniqueness. □

A.5 Option Value and Equilibrium Construction

Define the option-value function, using integration by parts:

$$\Phi(x) \equiv \mathbb{E}\left[(x-c)^+\right] = \int_{-\infty}^{x} G(t)\,dt, \tag{A.18}$$

where G is extended by zero below $\underline{c}$ and by one above $\bar{c}$. The value identity can then be written as

$$v(\gamma, N) = -\frac{a}{2\gamma} - \frac{a}{2\delta} + \Phi\big(c^*(\gamma, N)\big).$$

Proof of Theorem 2. Let

$$d(\eta) \equiv \frac{a}{2\eta} + \frac{a}{2\delta}.$$

The function Φ is continuous, increasing, convex, and 1-Lipschitz. Because waiting costs are strictly positive,

$$\Phi(x) < x \qquad \text{for every } x > 0. \tag{A.19}$$

Fix an arbitrary state (γ, N), let $m = N - 1$, and index the states along its resolution path by

$$(\gamma_k, N_k) \equiv (\gamma + (m-k)\delta, k+1)\,, \qquad k = 0, \ldots, m.$$

Thus state k contains bank i and k other unresolved banks. For notational convenience, write

$$d_k = d(\gamma_k), \qquad x_k = c^*(\gamma_k, N_k), \qquad q_k = 1 - G(x_k), \qquad v_k = v(\gamma_k, N_k).$$

At state k, the cutoff and value equations are

$$x_k = d_k + \mathbb{E}[v_{k-J}], \qquad J \sim \text{Binomial}(k, q_k), \tag{A.20}$$

$$v_k = -d_k + \Phi(x_k). \tag{A.21}$$

The event $J = 0$ returns the bank to the current state and therefore contains v_k itself.

Induction on k establishes both uniqueness of the cutoff and value at each state and the continuation-value ordering

$$v_0 > v_1 > \cdots > v_k.$$

When $k = 0$, equations (A.20)–(A.21) imply

$$x_0 = d_0 + v_0 = \Phi(x_0).$$

By (A.19), the unique solution is

$$x_0 = 0, \qquad q_0 = 1, \qquad v_0 = -d_0.$$

Now suppose that the equilibrium at states $0, \ldots, k-1$ is uniquely determined and that

$$v_0 > v_1 > \cdots > v_{k-1}.$$

Consider a candidate cutoff x at state k, and let

$$r = G(x), \qquad q = 1 - r.$$

If no peer resolves, which occurs with probability r^k, the candidate current value is $-d_k + \Phi(x)$. Substituting this value into the $J = 0$ term of (A.20), the cutoff equation becomes

$$x = T_k(x), \tag{A.22}$$

where

$$T_k(x) \equiv r^k \Phi(x) + \sum_{j=1}^{k} \binom{k}{j} q^j r^{k-j} K_j, \qquad K_j \equiv d_k + v_{k-j}. \tag{A.23}$$

The induction hypothesis implies

$$K_1 < K_2 < \cdots < K_k. \tag{A.24}$$

We first establish uniqueness. Any root of (A.22) must lie in $(0, \bar{c})$. Indeed, if $x \leq 0$, then $q = 1$ and

$$T_k(x) = K_k = d_k + v_0 = d_k - d_0 > 0,$$

so $x = T_k(x)$ is impossible. If $x \geq \bar{c}$, then $q = 0$ and the fixed-point equation reduces to $x = \Phi(x)$, contradicting (A.19).

For $q > 0$, define

$$C_k(q) \equiv \mathbb{E}[K_J \mid J \geq 1], \qquad J \sim \text{Binomial}(k, q).$$

Equation (A.22) is equivalent to

$$x - \Phi(x) = \left(1 - G(x)^k\right) \left[C_k(1 - G(x)) - \Phi(x)\right]. \tag{A.25}$$

The left-hand side is strictly increasing on $(0, \bar{c})$, since for $0 < x < y < \bar{c}$,

$$[y - \Phi(y)] - [x - \Phi(x)] = \int_x^y [1 - G(t)]\, dt > 0.$$

By (A.24), $C_k(q)$ is nondecreasing in q. To see this, conditional on $J \geq 1$, the binomial probabilities are proportional to

$$\binom{k}{j} \left(\frac{q}{1-q}\right)^j, \qquad j = 1, \ldots, k,$$

which has the monotone-likelihood-ratio property in q. Consequently, as x increases, both

$$1 - G(x)^k \quad \text{and} \quad C_k(1 - G(x)) - \Phi(x)$$

are nonincreasing. At any root of (A.25), the second expression is positive because $x - \Phi(x) > 0$. Hence, if $x < y$ were two roots, the left-hand side would be strictly larger at y, while the right-hand side would be weakly smaller. This is impossible. Therefore, there is at most one equilibrium cutoff at state k.

It remains to establish existence and the continuation-value ordering. Define $\bar{x}_k$ as the unique

solution to

$$\Phi(\bar{x}_k) = d_k + v_{k-1}. \tag{A.26}$$

Such a solution exists because

$$d_k + v_{k-1} = d_k - d_{k-1} + \Phi(x_{k-1}) > 0,$$

and Φ is strictly increasing above $\underline{c}$ and unbounded. Equation (A.26) also gives

$$\Phi(\bar{x}_k) - \Phi(x_{k-1}) = d_k - d_{k-1} \equiv D_k > 0. \tag{A.27}$$

Thus $\bar{x}_k > x_{k-1}$. Moreover, because Φ is 1-Lipschitz and $x_{k-1} < \bar{c}$,

$$\bar{x}_k - x_{k-1} > D_k. \tag{A.28}$$

It follows that

$$\bar{q}_k \equiv 1 - G(\bar{x}_k) \leq 1 - G(x_{k-1}) = q_{k-1}.$$

At the candidate cutoff $\bar{x}_k$, the candidate current value is

$$-d_k + \Phi(\bar{x}_k) = v_{k-1}.$$

Define

$$\widetilde{v}_k = v_{k-1}, \qquad \widetilde{v}_r = v_r \quad \text{for } r = 0, \ldots, k-1.$$

Let

$$\bar{J} \sim \text{Binomial}(k, \bar{q}_k), \qquad L \sim \text{Binomial}(k-1, q_{k-1}).$$

Since $\bar{q}_k \leq q_{k-1}$, these random variables can be coupled so that

$$\bar{J} \leq 1 + L \qquad \text{almost surely.}$$

The induction hypothesis implies that

$$\widetilde{v}_k,\ v_{k-1},\ v_{k-2}, \ldots, v_0$$

is nondecreasing in the number of peer resolutions. Therefore,

$$\mathbb{E}\left[\widetilde{v}_{k-\bar{J}}\right] \leq \mathbb{E}\left[v_{k-1-L}\right].$$

Using the cutoff equation at state $k-1$,

$$x_{k-1} = d_{k-1} + \mathbb{E}\left[v_{k-1-L}\right],$$

we obtain

$$\begin{aligned} T_k(\bar{x}_k) &= d_k + \mathbb{E}\left[\widetilde{v}_{k-\bar{J}}\right] \\ &\leq d_k + \mathbb{E}\left[v_{k-1-L}\right] \\ &= x_{k-1} + D_k < \bar{x}_k, \end{aligned} \tag{A.29}$$

where the final inequality follows from (A.28). At the other endpoint,

$$T_k(0) = d_k + v_0 = d_k - d_0 > 0. \tag{A.30}$$

Continuity of T_k therefore gives a root

$$x_k \in (0, \bar{x}_k).$$

The preceding monotone-crossing argument implies that this root is unique. Finally,

$$v_k = -d_k + \Phi(x_k) < -d_k + \Phi(\bar{x}_k) = v_{k-1}.$$

This establishes existence, uniqueness, and the strict value ordering at state k, completing the induction.

Because the initial state (γ, N) was arbitrary, the argument applies throughout the state space. Moving from k to $k-1$ gives

$$v(\gamma + \delta, N - 1) > v(\gamma, N).$$

Finally, $x_k < \bar{c}$ implies $q_k > 0$ at every active state. Since the number of unresolved banks is finite, the resulting recursive equilibrium is well defined and reaches absorption almost surely. □

A.6 Proofs of the Strategic-Waiting Results

Proof of Theorem 3. Fix (γ, N), let $k = N - 1$, and define

$$h_n \equiv v(\gamma + n\delta, N - n), \qquad n = 0, \ldots, k.$$

By Theorem 2, a realized peer resolution strictly raises the continuation value of a bank that remains unresolved. Therefore,

$$h_{n+1} > h_n, \qquad n = 0, \ldots, k-1. \tag{A.31}$$

Suppose that each of the k peer banks resolves in the current period with probability q, while continuation play from the next period onward follows the equilibrium characterized in Theorem 2. Bank i's cutoff is

$$c_i^*(\gamma, N; q) = \frac{a}{2\gamma} + \frac{a}{2\delta} + \mathbb{E}\left[h_J\right], \qquad J \sim \text{Binomial}(k, q). \tag{A.32}$$

Let $L \sim \text{Binomial}(k-1, q)$. Differentiating the binomial expectation in (A.32) gives

$$\frac{\partial c_i^*(\gamma, N; q)}{\partial q} = k\,\mathbb{E}\left[h_{L+1} - h_L\right] > 0, \tag{A.33}$$

where the strict inequality follows from (A.31). At $q = 0$ and $q = 1$, the same conclusion holds for the corresponding one-sided derivatives.

Equivalently, bank i's payoff advantage from resolving rather than waiting is

$$A_i(c, q) = c - c_i^*(\gamma, N; q).$$

It follows that

$$\frac{\partial A_i(c,q)}{\partial q} = -\frac{\partial c_i^*(\gamma, N; q)}{\partial q} < 0.$$

Thus current resolution decisions are strict strategic substitutes.

Finally, letting $g = G'$, whenever the cutoff lies in the interior of the support,

$$\frac{\partial q_i^{\mathrm{BR}}(\gamma, N; q)}{\partial q} = -g\big(c_i^*(\gamma, N; q)\big)\frac{\partial c_i^*(\gamma, N; q)}{\partial q} < 0, \tag{A.34}$$

because $g > 0$ on $(\underline{c}, \bar{c})$. Outside the support, G is locally constant, so the best-response probability is only weakly decreasing.

At an equilibrium state with delay, $q^*(\gamma, N) < 1$, which implies $c^*(\gamma, N) > \underline{c}$. Theorem 2 also gives $c^*(\gamma, N) < \bar{c}$. Hence the equilibrium cutoff is in the interior of the support, and the best-response probability is strictly decreasing at that equilibrium. Because extension is the complementary action, its best-response probability is weakly increasing in peer resolution and is strictly increasing at every equilibrium state with delay. □

A.7 Proof of Across-Period Complementarity

Proof of Theorem 4. Hold (δ, a) fixed and suppress them from the notation. Let

$$K = N_t - 1,$$

and index the states on the realized-resolution path by

$$(\gamma_k, N_k) = (\gamma_t + (K-k)\delta, k+1), \qquad k = 0, \ldots, K.$$

At state k, bank i and k peers remain unresolved. The main text indexes this path by cumulative resolutions n, whereas this proof uses remaining peers k. Since $n = K - k$,

$$c_n^* = x_{K-n}, \qquad q_n^{t+1} = q_{K-n}, \tag{A.35}$$

where

$$d_k \equiv \frac{a}{2\gamma_k} + \frac{a}{2\delta}, \qquad x_k \equiv c^*(\gamma_k, N_k), \qquad q_k \equiv 1 - G(x_k), \qquad v_k \equiv v(\gamma_k, N_k).$$

By Theorem 2,

$$v_{k-1} > v_k, \qquad k = 1, \ldots, K. \tag{A.36}$$

Define the reduction in the uncertainty discount produced by one additional resolution at state k by

$$D_k \equiv d_k - d_{k-1} = \frac{a\delta}{2\gamma_k(\gamma_k + \delta)} > 0.$$

Because γ_k decreases with k,

$$0 < D_1 \le D_2 \le \cdots \le D_K.$$

Thus, the value of the precision generated by an additional resolution is greater at less-informed

states.

We prove by induction that

$$0 = x_0 < x_1 < \cdots < x_K. \tag{A.37}$$

The terminal-state result $x_0 = 0$ gives the base case. Now suppose that

$$x_0 < x_1 < \cdots < x_{k-1}$$

for some $k \geq 1$.

For a candidate cutoff x at state k, let $T_k(x)$ denote the right-hand side of the scalar cutoff equation:

$$T_k(x) = G(x)^k \Phi(x) + \sum_{j=1}^{k} \binom{k}{j} [1 - G(x)]^j G(x)^{k-j} \big(d_k + v_{k-j}\big). \tag{A.38}$$

The equilibrium cutoff x_k is the unique solution to

$$x = T_k(x).$$

Evaluate the map at x_{k-1}. Define the candidate current value

$$\widehat{v}_k \equiv -d_k + \Phi(x_{k-1}),$$

and set

$$\widehat{v}_r \equiv v_r, \qquad r = 0, \ldots, k-1.$$

Let

$$L \sim \text{Binomial}(k-1, q_{k-1}).$$

At the candidate cutoff x_{k-1}, the k peer-resolution draws can be decomposed into the $k-1$ draws summarized by L and one additional independent Bernoulli draw with success probability q_{k-1}. Therefore,

$$\begin{aligned} T_k(x_{k-1}) = d_k &+ q_{k-1}\, \mathbb{E}\big[v_{k-1-L}\big] \\ &+ (1 - q_{k-1})\, \mathbb{E}\big[\widehat{v}_{k-L}\big]. \end{aligned} \tag{A.39}$$

At state $k-1$, the equilibrium cutoff satisfies

$$x_{k-1} = d_{k-1} + \mathbb{E}\big[v_{k-1-L}\big].$$

Subtracting this equation from (A.39) gives

$$T_k(x_{k-1}) - x_{k-1} = D_k + (1 - q_{k-1}) \mathbb{E}\left[\widehat{v}_{k-L} - v_{k-1-L}\right]. \tag{A.40}$$

If $L = 0$, then

$$\widehat{v}_k - v_{k-1} = -d_k + \Phi(x_{k-1}) - [-d_{k-1} + \Phi(x_{k-1})] = -D_k.$$

If $L \geq 1$, let $\ell = k - L$. Then $1 \leq \ell \leq k-1$, and the value identity gives

$$\begin{aligned} 0 < v_{\ell-1} - v_\ell &= D_\ell - [\Phi(x_\ell) - \Phi(x_{\ell-1})] \\ &\leq D_\ell \leq D_k. \end{aligned}$$

The strict positivity follows from (A.36), while the first inequality uses the induction hypothesis and the monotonicity of Φ. Consequently, for every realization of L,

$$\widehat{v}_{k-L} - v_{k-1-L} \geq -D_k.$$

Substituting this bound into (A.40) yields

$$T_k(x_{k-1}) - x_{k-1} \geq q_{k-1} D_k > 0. \tag{A.41}$$

The strict inequality follows from Theorem 2, which gives $q_{k-1} > 0$, and from $D_k > 0$.

For $x \geq \bar{c}$, no peer resolves, and therefore

$$T_k(x) - x = \Phi(x) - x < 0.$$

Continuity implies that the state-k fixed-point equation has a root strictly above x_{k-1}. By Theorem 2, this root is unique. Hence

$$x_k > x_{k-1},$$

which completes the induction and proves (A.37).

Finally, a state following n resolutions corresponds to $k = K - n$. Using (A.35),

$$c^*_{n+1} = x_{K-n-1} < x_{K-n} = c^*_n.$$

Since G is nondecreasing,

$$q^{t+1}_{n+1} = 1 - G(c^*_{n+1}) \geq 1 - G(c^*_n) = q^{t+1}_n.$$

Therefore,

$$\Delta_n q_{t+1}(\gamma_t, N_t; \delta) = q^{t+1}_{n+1} - q^{t+1}_n \geq 0,$$

which proves (3.11).

If $q^{t+1}_n < 1$, then $c^*_n > \underline{c}$. Because $c^*_{n+1} < c^*_n$ and G has strictly positive density on its support,

$$G(c^*_{n+1}) < G(c^*_n).$$

It follows that

$$q^{t+1}_{n+1} > q^{t+1}_n, \qquad \Delta_n q_{t+1}(\gamma_t, N_t; \delta) > 0,$$

which proves (3.12). □

A.8 Baseline Comparative Statics

Proposition 10. Baseline comparative statics of equilibrium resolution and delay.

1. *Number of unresolved banks.* Fix (γ, δ, a). For every $N \in \{1, \dots, \bar{N} - 1\}$,

$$c^*(\gamma, N+1; \delta, a) \geq c^*(\gamma, N; \delta, a), \tag{A.42}$$
$$q^*(\gamma, N+1; \delta, a) \leq q^*(\gamma, N; \delta, a). \tag{A.43}$$

   Thus, the presence of more unresolved peers increases an individual bank's equilibrium incentive to wait.

2. *Current public precision.* Fix (N, δ, a). If $\gamma^H > \gamma^L$, then

$$c^*(\gamma^H, N; \delta, a) \leq c^*(\gamma^L, N; \delta, a), \tag{A.44}$$
$$q^*(\gamma^H, N; \delta, a) \geq q^*(\gamma^L, N; \delta, a). \tag{A.45}$$

   Consequently,

$$D^{\text{ind}}(\gamma^H, N; \delta, a) \leq D^{\text{ind}}(\gamma^L, N; \delta, a), \tag{A.46}$$
$$D^{\text{all}}(\gamma^H, N; \delta, a) \leq D^{\text{all}}(\gamma^L, N; \delta, a). \tag{A.47}$$

   Thus, greater information already incorporated into the public state accelerates equilibrium resolution.

3. *Anticipated signal informativeness.* Fix (γ, N, a). If $\delta^H > \delta^L$, then

$$c^*(\gamma, N; \delta^H, a) \geq c^*(\gamma, N; \delta^L, a), \tag{A.48}$$
$$q^*(\gamma, N; \delta^H, a) \leq q^*(\gamma, N; \delta^L, a). \tag{A.49}$$

   Thus, more informative potential future resolution signals increase the current incentive to wait.

4. *Developer risk aversion.* Fix (γ, N, δ). If $a^H > a^L$, then

$$c^*(\gamma, N; \delta, a^H) \geq c^*(\gamma, N; \delta, a^L), \tag{A.50}$$
$$q^*(\gamma, N; \delta, a^H) \leq q^*(\gamma, N; \delta, a^L). \tag{A.51}$$

   Consequently,

$$D^{\text{ind}}(\gamma, N; \delta, a^H) \geq D^{\text{ind}}(\gamma, N; \delta, a^L), \tag{A.52}$$
$$D^{\text{all}}(\gamma, N; \delta, a^H) \geq D^{\text{all}}(\gamma, N; \delta, a^L). \tag{A.53}$$

   Thus, greater developer risk aversion increases equilibrium waiting and lengthens the resolution process.

Parts 1 and 2 vary the reduced state, N and γ; Parts 3 and 4 vary the primitives, signal precision δ and developer risk aversion a. More peers lower each bank's resolution probability but increase

the number of potential information producers. More informative signals encourage waiting before they arrive but yield larger precision gains when resolutions occur. These opposing effects leave the impact of N and δ on expected resolution delay ambiguous.

Figure 4 illustrates the effect of initial public precision on resolution delay. Appendix A.11 illustrates the current resolution-probability comparisons for N and δ, their effects on total delay, and sensitivity to the waiting-cost distribution.

Proof of Proposition 10. Throughout the proof, the waiting-cost distribution G is held fixed. Write

$$d(\gamma;\delta,a) \equiv \frac{a}{2\gamma} + \frac{a}{2\delta}, \qquad x(\gamma,N;\delta,a) \equiv c^*(\gamma,N;\delta,a).$$

For a candidate cutoff x, let $r = G(x)$ and $q = 1-r$. After substituting the candidate current value $-d(\gamma;\delta,a)+\Phi(x)$ into the event in which no peer resolves, the cutoff equation at a state with $N \geq 2$ can be written as

$$x = T^{\delta,a}_{\gamma,N}(x), \tag{A.54}$$

where

$$T^{\delta,a}_{\gamma,N}(x) = r^{N-1}\Phi(x) + \sum_{n=1}^{N-1}\binom{N-1}{n} q^n r^{N-1-n}\left[d(\gamma;\delta,a) + v(\gamma+n\delta, N-n;\delta,a)\right]. \tag{A.55}$$

The proof of Theorem 2 establishes that (A.54) has a unique solution. It also establishes the following single-crossing property:

$$T^{\delta,a}_{\gamma,N}(x) - x \begin{cases} > 0, & x < c^*(\gamma,N;\delta,a), \\ < 0, & x > c^*(\gamma,N;\delta,a). \end{cases}$$

We use this property repeatedly to translate an ordering of fixed-point maps into an ordering of equilibrium cutoffs. Since G is nondecreasing, each cutoff ordering implies the reverse ordering of $q^* = 1 - G(x)$.

Part 1: Number of unresolved banks.

Fix (γ,δ,a), and suppress these arguments where no confusion arises. We prove by induction on N that

$$x(\gamma,N+1;\delta,a) \geq x(\gamma,N;\delta,a), \tag{A.56}$$
$$v(\gamma,N+1;\delta,a) \geq v(\gamma,N;\delta,a). \tag{A.57}$$

The argument also uses the fact that equilibrium value is nondecreasing in current public precision:

$$v(\eta^H,M;\delta,a) \geq v(\eta^L,M;\delta,a) \qquad \text{whenever } \eta^H > \eta^L. \tag{A.58}$$

Part 2 establishes this auxiliary result independently of the market-size comparison.

Suppose that the market-size orderings hold for all comparisons involving at most N unresolved

banks. Let

$$x_N = x(\gamma, N; \delta, a), \qquad q_N = 1 - G(x_N),$$

and consider the cutoff map for a state with $N+1$ unresolved banks, evaluated at x_N. Couple the peer-resolution counts by writing

$$J \sim \text{Binomial}(N-1, q_N), \qquad B \sim \text{Bernoulli}(q_N),$$

independently. Thus, J is the number of peer resolutions in the N-bank state, while $J+B$ is the corresponding number in the $N+1$-bank state at the candidate cutoff x_N.

If $B = 0$ and $J = 0$, the candidate continuation value in the $N+1$-bank state is

$$-d(\gamma; \delta, a) + \Phi(x_N) = v(\gamma, N; \delta, a).$$

If $B = 0$ and $J \geq 1$, the induction hypothesis gives

$$v(\gamma + J\delta, N+1-J; \delta, a) \geq v(\gamma + J\delta, N-J; \delta, a).$$

If $B = 1$, the two successor states contain the same number of unresolved banks, but the successor state in the larger economy has one additional signal. By (A.58),

$$v(\gamma + (J+1)\delta, N-J; \delta, a) \geq v(\gamma + J\delta, N-J; \delta, a).$$

The candidate continuation value in the $N+1$-bank state is therefore weakly higher in every realization of (J, B). It follows that

$$T^{\delta,a}_{\gamma,N+1}(x_N) \geq T^{\delta,a}_{\gamma,N}(x_N) = x_N.$$

The single-crossing property implies

$$x(\gamma, N+1; \delta, a) \geq x_N.$$

Using the value identity,

$$\begin{aligned} v(\gamma, N+1; \delta, a) &= -d(\gamma; \delta, a) + \Phi\big(x(\gamma, N+1; \delta, a)\big) \\ &\geq -d(\gamma; \delta, a) + \Phi(x_N) = v(\gamma, N; \delta, a), \end{aligned}$$

which closes the induction. The argument includes the base comparison between $N = 1$ and $N = 2$, since

$$x(\gamma, 1; \delta, a) = 0, \qquad v(\gamma, 1; \delta, a) = -d(\gamma; \delta, a).$$

This proves Part 1.

Part 2: Current public precision.

Fix (N, δ, a). We first prove by induction on N that the equilibrium cutoff is nonincreasing in current public precision. When $N = 1$,

$$x(\gamma, 1; \delta, a) = 0$$

for every γ.

Now let $N \geq 2$, and suppose that the result holds at every state with fewer than N unresolved banks. For $n \geq 1$, define

$$\begin{aligned} K_n(\gamma) &\equiv d(\gamma;\delta,a) + v(\gamma + n\delta, N-n;\delta,a) \\ &= d(\gamma;\delta,a) - d(\gamma+n\delta;\delta,a) + \Phi\big(x(\gamma+n\delta, N-n;\delta,a)\big). \end{aligned} \tag{A.59}$$

The direct gain from the precision generated by n resolutions is strictly decreasing in current precision:

$$\frac{\partial}{\partial\gamma}\left[d(\gamma;\delta,a) - d(\gamma+n\delta;\delta,a)\right] = -\frac{a}{2\gamma^2} + \frac{a}{2(\gamma+n\delta)^2} < 0. \tag{A.60}$$

By the induction hypothesis,

$$x(\gamma + n\delta, N-n;\delta,a)$$

is nonincreasing in γ. Since Φ is increasing, (A.59) implies

$$K_n(\gamma^H) \leq K_n(\gamma^L) \qquad \text{whenever } \gamma^H > \gamma^L.$$

The weights in (A.55) are the same at a given candidate cutoff under the two current precisions. Hence

$$T^{\delta,a}_{\gamma^H,N}(x) \leq T^{\delta,a}_{\gamma^L,N}(x) \qquad \text{for every } x.$$

Evaluating this inequality at the lower-precision equilibrium cutoff and applying the single-crossing property gives

$$x(\gamma^H, N;\delta,a) \leq x(\gamma^L, N;\delta,a). \tag{A.61}$$

We next establish the auxiliary value ordering (A.58). The result holds at $N = 1$ because

$$v(\gamma, 1;\delta,a) = -d(\gamma;\delta,a)$$

is increasing in γ. Suppose it holds at all states with fewer than N unresolved banks. Write

$$q^\sigma = q^*(\gamma^\sigma, N;\delta,a), \qquad r^\sigma = 1 - q^\sigma, \qquad \sigma \in \{H, L\}.$$

For a candidate current value w, define

$$\begin{aligned} \mathcal{B}^\sigma(w) \equiv \mathbb{E}_c\Bigg[\max\Bigg\{ &-d(\gamma^\sigma;\delta,a), -c + (r^\sigma)^{N-1}w \\ &+ \sum_{n=1}^{N-1}\binom{N-1}{n}(q^\sigma)^n(r^\sigma)^{N-1-n}v(\gamma^\sigma + n\delta, N-n;\delta,a)\Bigg\}\Bigg]. \end{aligned}$$

The equilibrium value $v(\gamma^\sigma, N;\delta,a)$ is the unique fixed point of $\mathcal{B}^\sigma$. The operator is increasing and has contraction modulus at most $(r^\sigma)^{N-1} < 1$.

Set

$$w = v(\gamma^L, N;\delta,a).$$

Under the lower-precision state, the sequence

$$w, \quad v(\gamma^L + \delta, N-1; \delta, a), \quad \ldots, \quad v(\gamma^L + (N-1)\delta, 1; \delta, a)$$

is increasing by Theorem 2. At every successor state with fewer than N unresolved banks, the induction hypothesis gives

$$v(\gamma^H + n\delta, N-n; \delta, a) \geq v(\gamma^L + n\delta, N-n; \delta, a).$$

Moreover, (A.61) implies $q^H \geq q^L$, so the peer-resolution count under the higher-precision state first-order stochastically dominates that under the lower-precision state. Together with

$$-d(\gamma^H; \delta, a) > -d(\gamma^L; \delta, a),$$

these observations imply

$$\mathcal{B}^H(w) \geq \mathcal{B}^L(w) = w.$$

Iterating the increasing contraction $\mathcal{B}^H$ from w gives

$$v(\gamma^H, N; \delta, a) \geq v(\gamma^L, N; \delta, a),$$

which completes the induction and proves (A.58).

It remains to establish the delay orderings. Consider two economies with the same initial set of banks, the same (δ, a), and initial precisions $\gamma^H > \gamma^L$. Couple the economies by assigning every bank the same sequence of waiting-cost draws.

We claim that the set of unresolved banks in the higher-precision economy is a subset of the corresponding set in the lower-precision economy at every date. Suppose the inclusion holds at the beginning of a period, and let $N_t^H \leq N_t^L$ be the corresponding numbers of unresolved banks. Their precisions satisfy

$$\gamma_t^H = \gamma_t^L + (N_t^L - N_t^H)\delta + (\gamma^H - \gamma^L).$$

Repeated application of Theorem 4 gives

$$c^*\big(\gamma_t^L + (N_t^L - N_t^H)\delta, N_t^H; \delta, a\big) \leq c^*(\gamma_t^L, N_t^L; \delta, a).$$

Combining this inequality with (A.61) yields

$$c^*(\gamma_t^H, N_t^H; \delta, a) \leq c^*(\gamma_t^L, N_t^L; \delta, a).$$

With common cost draws, any bank present in both economies that resolves in the lower-precision economy also resolves in the higher-precision economy. The set inclusion is therefore preserved.

Consequently, every bank resolves weakly earlier in the higher-precision economy, and the last bank also resolves weakly earlier. Taking expectations gives

$$D^{\text{ind}}(\gamma^H, N; \delta, a) \leq D^{\text{ind}}(\gamma^L, N; \delta, a),$$
$$D^{\text{all}}(\gamma^H, N; \delta, a) \leq D^{\text{all}}(\gamma^L, N; \delta, a).$$

This proves Part 2.

Part 3: Anticipated signal informativeness.

Fix (γ, N, a). For this comparative-technology argument only, define the auxiliary shifted value

$$w(\gamma, N; \delta, a) \equiv v(\gamma, N; \delta, a) + \frac{a}{2\delta}, \qquad d_0(\gamma) \equiv \frac{a}{2\gamma}.$$

This leaves the main text's definition of v unchanged. The identity $d(\gamma; \delta, a) + v(\eta, M; \delta, a) = d_0(\gamma) + w(\eta, M; \delta, a)$ removes the constant discount from the cutoff equation. At a fixed technology, the continuation-value and public-precision orderings for v apply equally to w. We prove jointly by induction on N that the cutoff and shifted value are nondecreasing in δ. When $N = 1$, there are no peer-generated signals, and

$$x(\gamma, 1; \delta, a) = 0, \qquad w(\gamma, 1; \delta, a) = -d_0(\gamma),$$

independently of δ.

Now let $N \geq 2$, suppose that the result holds at every state with fewer than N unresolved banks, and take $\delta^H > \delta^L$. For $n = 1, \ldots, N-1$, define

$$V_n^H \equiv w(\gamma + n\delta^H, N - n; \delta^H, a),$$
$$V_n^L \equiv w(\gamma + n\delta^L, N - n; \delta^L, a).$$

By (A.58),

$$w(\gamma + n\delta^H, N - n; \delta^H, a) \geq w(\gamma + n\delta^L, N - n; \delta^H, a).$$

The induction hypothesis, applied at the fixed reduced state $(\gamma + n\delta^L, N - n)$, gives

$$w(\gamma + n\delta^L, N - n; \delta^H, a) \geq w(\gamma + n\delta^L, N - n; \delta^L, a).$$

Therefore,

$$V_n^H \geq V_n^L, \qquad n = 1, \ldots, N-1. \tag{A.62}$$

At a given candidate cutoff, the current uncertainty-discount term $d_0(\gamma)$ and the binomial weights are the same under the two signal technologies. Equation (A.62) therefore implies

$$T_{\gamma,N}^{\delta^H,a}(x) \geq T_{\gamma,N}^{\delta^L,a}(x) \qquad \text{for every } x.$$

Evaluating at the equilibrium cutoff under δ^L and applying the single-crossing property gives

$$x(\gamma, N; \delta^H, a) \geq x(\gamma, N; \delta^L, a). \tag{A.63}$$

The shifted value identity is

$$w(\gamma, N; \delta, a) = -d_0(\gamma) + \Phi\big(x(\gamma, N; \delta, a)\big).$$

Thus, (A.63) also gives

$$w(\gamma, N; \delta^H, a) \geq w(\gamma, N; \delta^L, a),$$

closing the induction.

This proves Part 3.

Part 4: Developer risk aversion.

Fix (γ, N, δ). We prove by induction on N that the equilibrium cutoff is nondecreasing in a. When $N = 1$,

$$x(\gamma, 1; \delta, a) = 0$$

for every a.

Now let $N \geq 2$, and suppose that the result holds at every state with fewer than N unresolved banks. For $n \geq 1$, define

$$\begin{aligned} K_n(a) &\equiv d(\gamma; \delta, a) + v(\gamma + n\delta, N - n; \delta, a) \\ &= \frac{a}{2}\left(\frac{1}{\gamma} - \frac{1}{\gamma + n\delta}\right) + \Phi\big(x(\gamma + n\delta, N - n; \delta, a)\big). \end{aligned} \tag{A.64}$$

The first term in (A.64) is strictly increasing in a. The second term is nondecreasing in a by the induction hypothesis and the monotonicity of Φ. Hence, whenever $a^H > a^L$,

$$K_n(a^H) \geq K_n(a^L) \qquad \text{for every } n \geq 1.$$

At a given candidate cutoff, the binomial weights are independent of a. It follows that

$$T_{\gamma,N}^{\delta,a^H}(x) \geq T_{\gamma,N}^{\delta,a^L}(x) \qquad \text{for every } x.$$

Evaluating at the low-a equilibrium cutoff and applying the single-crossing property gives

$$x(\gamma, N; \delta, a^H) \geq x(\gamma, N; \delta, a^L). \tag{A.65}$$

To establish the delay orderings, consider two economies with the same initial reduced state and signal technology, and assign every bank the same sequence of waiting-cost draws. Let H denote the economy with $a^H > a^L$. We claim that the set of unresolved banks in the low-a economy is a subset of that in the high-a economy at every date.

Suppose the inclusion holds at the beginning of a period, and let $N_t^H \geq N_t^L$ denote the corresponding numbers of unresolved banks. Since both economies began with the same public precision,

$$\gamma_t^L = \gamma_t^H + (N_t^H - N_t^L)\delta.$$

The fixed-state ordering (A.65) gives

$$c^*(\gamma_t^H, N_t^H; \delta, a^H) \geq c^*(\gamma_t^H, N_t^H; \delta, a^L).$$

Repeated application of Theorem 4 in the low-a economy gives

$$c^*(\gamma_t^H, N_t^H; \delta, a^L) \geq c^*(\gamma_t^L, N_t^L; \delta, a^L).$$

Therefore,

$$c^*(\gamma_t^H, N_t^H; \delta, a^H) \geq c^*(\gamma_t^L, N_t^L; \delta, a^L).$$

Any bank present in both economies that resolves in the high-a economy also resolves in the low-a economy, preserving the set inclusion.

It follows that each bank, as well as the last unresolved bank, resolves weakly later in the high-a economy. Hence

$$D^{\text{ind}}(\gamma, N; \delta, a^H) \geq D^{\text{ind}}(\gamma, N; \delta, a^L),$$
$$D^{\text{all}}(\gamma, N; \delta, a^H) \geq D^{\text{all}}(\gamma, N; \delta, a^L).$$

This proves Part 4 and completes the proof. □

A.9 Proofs of the Amplification Results

A.9.1 Signal Informativeness at a Common Current State

Lemma 11. Fix $N \geq 2$, $\gamma > 0$, $a > 0$, and $\delta^H > \delta^L > 0$. If

$$\gamma^2 \leq 2\delta^H \delta^L, \tag{A.66}$$

then

$$\Delta_0 q_{t+1}(\gamma, N; \delta^H) \geq \Delta_0 q_{t+1}(\gamma, N; \delta^L).$$

Proof. Let

$$d(\eta; \delta) \equiv \frac{a}{2\eta} + \frac{a}{2\delta}.$$

We first compare the resolution probabilities before the additional resolution occurs. Because the current state (γ, N) is held fixed, the anticipated-signal-informativeness result in Proposition 10 implies

$$q^*(\gamma, N; \delta^H) \leq q^*(\gamma, N; \delta^L). \tag{A.67}$$

It remains to compare the probabilities after one resolution has occurred. For $e \in \{H, L\}$ and $m = 1, \ldots, N-1$, define the states along the post-resolution path by

$$\eta_m^e \equiv \gamma + (N - m)\delta^e, \tag{A.68}$$
$$x_m^e \equiv c^*(\eta_m^e, m; \delta^e). \tag{A.69}$$

In particular,

$$x_{N-1}^e = c^*(\gamma + \delta^e, N - 1; \delta^e)$$

is the cutoff immediately following the additional resolution.

We prove by induction on m that

$$x_m^H \leq x_m^L, \qquad m = 1, \ldots, N-1. \tag{A.70}$$

When $m = 1$, the only remaining bank resolves immediately, so

$$x_1^H = x_1^L = 0.$$

Now fix $m \geq 2$ and suppose that (A.70) holds at every state with fewer than m unresolved banks. Let

$$\ell \equiv N - m \geq 1.$$

For $j = 1, \ldots, m-1$, define the uncertainty-discount gain from j additional resolutions by

$$\begin{aligned} A_{\ell,j}(\delta) &\equiv d(\gamma + \ell\delta; \delta) - d(\gamma + (\ell + j)\delta; \delta) \\ &= \frac{aj\delta}{2(\gamma + \ell\delta)(\gamma + (\ell + j)\delta)}. \end{aligned} \tag{A.71}$$

Direct cross multiplication shows that

$$A_{\ell,j}(\delta^H) \leq A_{\ell,j}(\delta^L) \quad \Longleftrightarrow \quad \delta^H \delta^L \geq \frac{\gamma^2}{\ell(\ell + j)}. \tag{A.72}$$

Because $\ell \geq 1$ and $j \geq 1$,

$$\ell(\ell + j) \geq 2.$$

Condition (A.66) therefore implies

$$A_{\ell,j}(\delta^H) \leq A_{\ell,j}(\delta^L) \qquad \text{for every } j = 1, \ldots, m-1. \tag{A.73}$$

For a candidate cutoff x, write

$$r = G(x), \qquad p = 1 - r.$$

Under technology e, the scalar cutoff map at state (η_m^e, m) is

$$\begin{aligned} T_m^e(x) = r^{m-1}\Phi(x) \\ + \sum_{j=1}^{m-1} \binom{m-1}{j} p^j r^{m-1-j} K_{m,j}^e, \end{aligned} \tag{A.74}$$

where

$$\begin{aligned} K_{m,j}^e &\equiv d(\eta_m^e; \delta^e) + v(\eta_m^e + j\delta^e, m - j; \delta^e) \\ &= A_{\ell,j}(\delta^e) + \Phi(x_{m-j}^e). \end{aligned} \tag{A.75}$$

The induction hypothesis and the monotonicity of Φ imply

$$\Phi(x_{m-j}^H) \leq \Phi(x_{m-j}^L).$$

Together with (A.73), this gives

$$K_{m,j}^H \leq K_{m,j}^L$$

for every j. Hence

$$T_m^H(x) \le T_m^L(x) \qquad \text{for every } x. \tag{A.76}$$

Let x_m^L be the unique fixed point of T_m^L. Then

$$T_m^H(x_m^L) \le T_m^L(x_m^L) = x_m^L.$$

Moreover, $T_m^H(0) > 0$. Continuity and equilibrium uniqueness therefore imply

$$x_m^H \le x_m^L.$$

This completes the induction.

Applying the result at $m = N-1$ and using the monotonicity of G gives

$$q^*(\gamma + \delta^H, N-1; \delta^H) \ge q^*(\gamma + \delta^L, N-1; \delta^L). \tag{A.77}$$

Finally,

$$\begin{aligned} &\Delta_0 q_{t+1}(\gamma, N; \delta^H) - \Delta_0 q_{t+1}(\gamma, N; \delta^L) \\ &= \underbrace{q^*(\gamma + \delta^H, N-1; \delta^H) - q^*(\gamma + \delta^L, N-1; \delta^L)}_{\ge 0} \\ &\quad + \underbrace{q^*(\gamma, N; \delta^L) - q^*(\gamma, N; \delta^H)}_{\ge 0}, \end{aligned}$$

where the first inequality follows from (A.77) and the second from (A.67). Therefore,

$$\Delta_0 q_{t+1}(\gamma, N; \delta^H) \ge \Delta_0 q_{t+1}(\gamma, N; \delta^L),$$

which proves the lemma. □

A.9.2 Signal Informativeness Along Resolution Histories

Proof of Proposition 5. Fix (γ_t, N_t), with $N_t \ge 3$, and $\delta^H > \delta^L > 0$. Hold the cost distribution G fixed. A sufficient low-precision condition is $\gamma_t^2 \le 2\delta^H\delta^L$. Let $\mu_c = \mathbb{E}[c] > 0$. The rare-resolution approximation established in (A.87) below implies, for $r = 0, \dots, N_t - 2$ and $e \in \{H, L\}$,

$$a q_r^{t+1}(\gamma_t, N_t; \delta^e) \longrightarrow C_r^e \equiv \frac{2\mu_c(\gamma_t + r\delta^e)(\gamma_t + (r+1)\delta^e)}{\delta^e(N_t - r - 1)}. \tag{A.78}$$

Both probability terms in each marginal effect are covered by this approximation when $n \le N_t - 3$. Direct subtraction gives

$$\begin{aligned} L_n &\equiv \lim_{a\to\infty} a\left[\Delta_n q_{t+1}(\gamma_t, N_t; \delta^H) - \Delta_n q_{t+1}(\gamma_t, N_t; \delta^L)\right] \\ &= \frac{2\mu_c(\delta^H - \delta^L)}{(N_t - n - 2)(N_t - n - 1)}\left[(n+1)(2N_t - n - 2) - \frac{\gamma_t^2}{\delta^H\delta^L}\right]. \end{aligned} \tag{A.79}$$

For every $n = 0, \ldots, N_t - 3$,

$$(n+1)(2N_t - n - 2) \geq 2(N_t - 1) \geq 4, \qquad \frac{\gamma_t^2}{\delta^H \delta^L} \leq 2.$$

Hence $L_n > 0$. Since the index set is finite, convergence gives a finite common threshold a^* above which all the stated inequalities hold strictly.

For an explicit sufficient bound, write

$$\eta_r^e = \gamma_t + r\delta^e, \qquad k_r = N_t - r - 1, \qquad b_r^e = \frac{\delta^e}{2\eta_r^e(\eta_r^e + \delta^e)}.$$

Then $C_r^e = \mu_c/(k_r b_r^e)$. For $ab_r^e \geq 2\bar{c}$, the bound in (A.86) implies

$$\left| \frac{C_r^e}{aq_r^{t+1}(\gamma_t, N_t; \delta^e)} - 1 \right| \leq \frac{R_r^e}{a}, \qquad R_r^e \equiv \frac{2(k_r - 1)\mu_c + (k_r + 1)\bar{c}/k_r}{b_r^e}.$$

Take $b_{\min}$ to be the minimum of b_r^e over the stated states, $C_{\max}$ and $R_{\max}$ the corresponding maxima, and $L_{\min} = \min_{0 \leq n \leq N_t - 3} L_n > 0$. A sufficient threshold is

$$a^* = \max\left\{ \frac{2\bar{c}}{b_{\min}},\ 2R_{\max},\ \frac{16 C_{\max} R_{\max}}{L_{\min}} \right\}. \tag{A.80}$$

The first two restrictions yield $|aq_r^{t+1} - C_r^e| \leq 2C_{\max}R_{\max}/a$. There are four probability terms in the comparison, so the final restriction bounds their total error by $L_{\min}/2$. Therefore

$$a\left[\Delta_n q_{t+1}(\gamma_t, N_t; \delta^H) - \Delta_n q_{t+1}(\gamma_t, N_t; \delta^L)\right] \geq L_{\min}/2 > 0.$$

The threshold depends on the fixed initial state, signal pair, and cost distribution. The terminal comparison $n = N_t - 2$ is excluded because its one-bank successor has resolution probability one. □

A.9.3 Current Public Precision and Proportional Amplification

Proof of Proposition 6. First establish the comparison for zero versus one resolution from a common current state with $N \geq 3$. Let

$$\mu_c \equiv \mathbb{E}[c], \qquad \bar{c} \equiv \sup \operatorname{supp}(G),$$

and define the value of the precision generated by one resolution at public precision η by

$$B(\eta; \delta, a) \equiv d(\eta) - d(\eta + \delta) = \frac{a\delta}{2\eta(\eta + \delta)}. \tag{A.81}$$

We first establish a rare-resolution approximation. Consider an arbitrary active state (η, m) with

$m \geq 2$. To simplify notation, write

$$x = c^*(\eta, m; \delta, a), \qquad q = q^*(\eta, m; \delta, a), \qquad k = m - 1,$$

and let

$$J \sim \text{Binomial}(k, q).$$

For $j \geq 1$, define

$$\begin{aligned} B_j &\equiv d(\eta) - d(\eta + j\delta), \\ x_j &\equiv c^*(\eta + j\delta, m - j; \delta, a). \end{aligned}$$

The scalar cutoff equation is

$$x = (1-q)^k \Phi(x) + \sum_{j=1}^{k} \Pr(J = j) \left[B_j + \Phi(x_j) \right]. \tag{A.82}$$

Define

$$h(x) \equiv x - \Phi(x) = \mathbb{E}[\min\{c, x\}].$$

Because

$$1 - (1-q)^k = \Pr(J \geq 1),$$

equation (A.82) can be rearranged as

$$h(x) = \sum_{j=1}^{k} \Pr(J = j) B_j + \sum_{j=1}^{k} \Pr(J = j) \left[\Phi(x_j) - \Phi(x) \right]. \tag{A.83}$$

Let $B \equiv B(\eta; \delta, a) = B_1$. Since d is decreasing and convex,

$$B \leq B_j \leq jB.$$

If $B > \bar{c}$, the original cutoff equation implies

$$\Pr(J \geq 1) B \leq \mu_c + \Pr(J \geq 1) \bar{c}.$$

Consequently,

$$q \leq \Pr(J \geq 1) \leq \frac{\mu_c}{B - \bar{c}}. \tag{A.84}$$

Moreover,

$$0 \leq kqB - \sum_{j=1}^{k} \Pr(J = j) B_j \leq kqB(k-1)q.$$

Because $0 \leq \Phi(x) \leq \bar{c}$, (A.83) therefore gives

$$|h(x) - kqB| \leq kqB(k-1)q + kq\bar{c}. \tag{A.85}$$

Finally,

$$0 \le \mu_c - h(x) = \mathbb{E}[(c-x)^+] \le \bar{c}q.$$

Combining this inequality with (A.84) and (A.85) yields

$$\left| \frac{\mu_c}{kqB} - 1 \right| \le (k-1)\frac{\mu_c}{B-\bar{c}} + \frac{k+1}{k}\frac{\bar{c}}{B}. \tag{A.86}$$

Because $B(\eta;\delta,a)$ grows linearly in a, the right-hand side converges to zero as $a \to \infty$. Thus, for every fixed $m \ge 2$ and $\eta > 0$,

$$q^*(\eta, m; \delta, a) \sim \frac{\mu_c}{(m-1)B(\eta;\delta,a)} \qquad \text{as } a \to \infty. \tag{A.87}$$

Apply (A.87) to the pre-resolution state (γ, N) and the successor state $(\gamma+\delta, N-1)$. Since

$$\frac{B(\gamma;\delta,a)}{B(\gamma+\delta;\delta,a)} = \frac{\gamma+2\delta}{\gamma},$$

we obtain

$$\Delta_0 \log q_{t+1}(\gamma, N; \delta, a) \longrightarrow \log\left(\frac{N-1}{N-2}\right) + \log\left(1+\frac{2\delta}{\gamma}\right) \tag{A.88}$$

as $a \to \infty$. The first term is the proportional effect of having one fewer unresolved bank; the second is the proportional effect of the additional public precision.

It follows that

$$\begin{aligned} \lim_{a\to\infty} &\left[\Delta_0 \log q_{t+1}(\gamma^L, N; \delta, a) - \Delta_0 \log q_{t+1}(\gamma^H, N; \delta, a)\right] \\ &= \log\left(1+\frac{2\delta}{\gamma^L}\right) - \log\left(1+\frac{2\delta}{\gamma^H}\right) > 0. \end{aligned} \tag{A.89}$$

Therefore, convergence implies that there exists a finite a^* such that the desired strict inequality holds for every sufficiently large a.

An explicit sufficient threshold follows from (A.86). Define

$$D_\infty \equiv \log\left(1+\frac{2\delta}{\gamma^L}\right) - \log\left(1+\frac{2\delta}{\gamma^H}\right) > 0, \tag{A.90}$$

$$\rho^* \equiv 1 - \exp\left(-\frac{D_\infty}{8}\right), \tag{A.91}$$

$$A \equiv (N-2)\mu_c, \tag{A.92}$$

$$C \equiv \rho^*\bar{c} + A + 2\bar{c}, \tag{A.93}$$

$$B^* \equiv \frac{C + \sqrt{C^2 - 8\rho^*\bar{c}^2}}{2\rho^*}. \tag{A.94}$$

Set

$$a^* \equiv \frac{2(\gamma^H + \delta)(\gamma^H + 2\delta)}{\delta} B^*. \tag{A.95}$$

For $a \geq a^*$, the smallest value of B among the four states entering the comparison is

$$B_{\min}(a) = B(\gamma^H + \delta; \delta, a) \geq B^*.$$

Equation (A.86) then implies, uniformly across these four states,

$$\left| \frac{\mu_c}{(m-1)q^*(\eta, m; \delta, a)B(\eta; \delta, a)} - 1 \right| \leq \rho^*.$$

The corresponding approximation error in each log probability is at most

$$-\log(1 - \rho^*) = \frac{D_\infty}{8}.$$

The comparison contains four log probabilities, so the total approximation error is at most $D_\infty/2$. Combining this bound with (A.89) gives

$$\Delta_0 \log q_{t+1}(\gamma^L, N; \delta, a) - \Delta_0 \log q_{t+1}(\gamma^H, N; \delta, a) \geq \frac{D_\infty}{2} > 0.$$

This proves the common-current-state comparison.

For the historical statement, apply this argument separately at each $n = 0, \ldots, N_t - 3$, replacing

$$(\gamma^L, \gamma^H, N) \quad \text{by} \quad (\gamma^L + n\delta, \gamma^H + n\delta, N_t - n).$$

At least three banks remain in each pre-event state, and the precision ordering is preserved. Equivalently, the limiting historical effect is

$$\Delta_n \log q_{t+1}(\gamma_t, N_t; \delta) \longrightarrow \log\left(\frac{N_t - n - 1}{N_t - n - 2}\right) + \log\left(1 + \frac{2\delta}{\gamma_t + n\delta}\right),$$

which is strictly decreasing in γ_t. Let a_n^* be the finite threshold obtained from (A.95) with the corresponding substitution. Taking $a^* = \max_{0 \leq n \leq N_t - 3} a_n^*$ establishes every stated historical comparison jointly and proves the proposition. □

A.10 Proofs for the Fixed-Waiting-Cost Benchmark

A.10.1 Equilibrium with a Fixed Waiting Cost

Proof of Proposition 7. The result follows by induction on N. When $N = 1$, waiting produces no information and costs $\kappa > 0$, so immediate resolution is strictly optimal and

$$q^\kappa(\gamma, 1; \delta, a) = 1, \qquad v^\kappa(\gamma, 1) = -d(\gamma).$$

Now consider $N \geq 2$, and suppose that (A.8) holds at all successor states with fewer than N

unresolved banks. An equilibrium with $q = 0$ is impossible: if no bank ever resolves, waiting returns the bank to the same state after imposing the strictly positive cost κ, so resolving is a profitable deviation.

Hence resolution occurs with positive probability. If the equilibrium is interior, a bank must be indifferent between resolving and waiting. Because resolving yields $-d(\gamma)$, its equilibrium value is therefore $-d(\gamma)$. The same value obtains when $q = 1$, since the bank resolves with probability one. Thus all continuation values entering the waiting payoff take the form

$$v^{\kappa}(\gamma + j\delta, N - j) = -d(\gamma + j\delta).$$

The payoff advantage of waiting rather than resolving, when each peer resolves with probability q, is consequently

$$\begin{aligned} &-\kappa + \mathbb{E}\left[-d(\gamma + J\delta)\right] - \left[-d(\gamma)\right] \\ &\quad = \mathcal{I}_{\gamma,N}(q; \delta, a) - \kappa. \end{aligned} \tag{A.96}$$

The function $\mathcal{I}_{\gamma,N}(q; \delta, a)$ is continuous, satisfies

$$\mathcal{I}_{\gamma,N}(0; \delta, a) = 0,$$

and is strictly increasing in q. To see the last property, let

$$L \sim \text{Binomial}(N - 2, q).$$

Differentiating the binomial expectation gives

$$\frac{\partial \mathcal{I}_{\gamma,N}(q; \delta, a)}{\partial q} = (N - 1)\mathbb{E}\left[d(\gamma + L\delta) - d(\gamma + (L + 1)\delta)\right] > 0. \tag{A.97}$$

At $q = 1$,

$$\mathcal{I}_{\gamma,N}(1; \delta, a) = d(\gamma) - d\big(\gamma + (N - 1)\delta\big).$$

If this quantity is weakly below κ, resolving is optimal when all peers resolve, and $q = 1$ is the unique symmetric equilibrium probability. Otherwise, continuity and strict monotonicity imply that there is a unique $q \in (0, 1)$ satisfying

$$\mathcal{I}_{\gamma,N}(q; \delta, a) = \kappa.$$

At this probability, each bank is indifferent and can mix with resolution probability q. This proves (A.7) and completes the induction. □

A.10.2 Strategic Substitution and Across-Period Complementarity

Proof of Proposition 8. Part 1 follows immediately from (A.97). When peers resolve more frequently, waiting provides a greater expected informational benefit. Because resolving forgoes that benefit, current resolution decisions are strategic substitutes.

For Part 2, fix $q > 0$. Let

$$L \sim \text{Binomial}(N-2, q), \qquad Z \sim \text{Bernoulli}(q),$$

independently. Before the additional resolution, the number of peer resolutions is distributed as $L + Z$. After the additional resolution, there are $N - 2$ remaining peers and the current precision is $\gamma + \delta$.

For every realization of (L, Z),

$$\begin{aligned} &d(\gamma) - d\big(\gamma + (L+Z)\delta\big) \\ &\quad \geq d(\gamma+\delta) - d\big(\gamma + (L+1)\delta\big) \end{aligned} \tag{A.98}$$

whenever $Z = 0$. This follows from the convexity of d: a fixed precision increment is more valuable when current precision is lower. When $Z = 1$, the left-hand side additionally includes the strictly positive first increment

$$d(\gamma) - d(\gamma + \delta),$$

so the inequality is strict. Because $q > 0$, the event $Z = 1$ has positive probability. Taking expectations therefore gives

$$\mathcal{I}_{\gamma,N}(q;\delta,a) > \mathcal{I}_{\gamma+\delta,N-1}(q;\delta,a). \tag{A.99}$$

Let

$$q_0 \equiv q^\kappa(\gamma, N; \delta, a), \qquad q_1 \equiv q^\kappa(\gamma+\delta, N-1; \delta, a).$$

If $q_0 < 1$, equilibrium indifference and (A.99) imply

$$\mathcal{I}_{\gamma+\delta,N-1}(q_0;\delta,a) < \mathcal{I}_{\gamma,N}(q_0;\delta,a) = \kappa.$$

Since the post-resolution information benefit is strictly increasing in q, restoring indifference requires $q_1 > q_0$, unless the post-resolution equilibrium is at the corner $q_1 = 1$, in which case the same strict inequality holds.

If $q_0 = 1$, then

$$\kappa \geq \mathcal{I}_{\gamma,N}(1;\delta,a) > \mathcal{I}_{\gamma+\delta,N-1}(1;\delta,a),$$

so the post-resolution equilibrium also has $q_1 = 1$. Therefore, $q_1 \geq q_0$, with strict inequality whenever $q_0 < 1$. □

A.10.3 Amplification in the Fixed-Cost Benchmark

Proof of Proposition 9. For Part 1, increasing δ raises

$$d(\gamma) - d(\gamma + j\delta)$$

for every $j \geq 1$. Hence, at the pre-resolution state,

$$q^\kappa(\gamma, N; \delta^H, a) \leq q^\kappa(\gamma, N; \delta^L, a).$$

At the post-resolution state, the informational benefit from j additional resolutions is

$$A_j(\delta) \equiv d(\gamma+\delta) - d(\gamma+(j+1)\delta) = \frac{aj\delta}{2(\gamma+\delta)(\gamma+(j+1)\delta)}.$$

For $j \geq 1$, direct cross multiplication gives

$$A_j(\delta^H) \leq A_j(\delta^L) \quad \Longleftrightarrow \quad \delta^H\delta^L \geq \frac{\gamma^2}{j+1}.$$

Because $j+1 \geq 2$, the maintained condition implies this ordering for every $j \geq 1$. The post-resolution information-benefit function is therefore weakly lower under δ^H. Its equilibrium resolution probability must consequently be weakly higher:

$$q^\kappa(\gamma+\delta^H, N-1; \delta^H, a) \geq q^\kappa(\gamma+\delta^L, N-1; \delta^L, a).$$

Combining the pre- and post-resolution orderings proves (A.12).

For Part 2, consider any state (η, m) with $m \geq 2$, and set

$$k = m-1, \qquad B(\eta;\delta,a) \equiv d(\eta) - d(\eta+\delta) = \frac{a\delta}{2\eta(\eta+\delta)}.$$

For sufficiently large a, the equilibrium is interior and satisfies

$$\kappa = \mathbb{E}\left[d(\eta) - d(\eta + J\delta)\right], \qquad J \sim \text{Binomial}(k, q^\kappa).$$

Because

$$B(\eta;\delta,a) \leq d(\eta) - d(\eta+j\delta) \leq jB(\eta;\delta,a),$$

we have

$$(1-q^\kappa)^{k-1} \leq \frac{\kappa}{kq^\kappa B(\eta;\delta,a)} \leq 1.$$

Moreover,

$$q^\kappa \leq \frac{\kappa}{B(\eta;\delta,a)} \longrightarrow 0 \qquad \text{as } a \to \infty.$$

It follows that

$$q^\kappa(\eta, m; \delta, a) \sim \frac{\kappa}{(m-1)B(\eta;\delta,a)}. \tag{A.100}$$

Applying (A.100) to (γ, N) and $(\gamma+\delta, N-1)$, and using

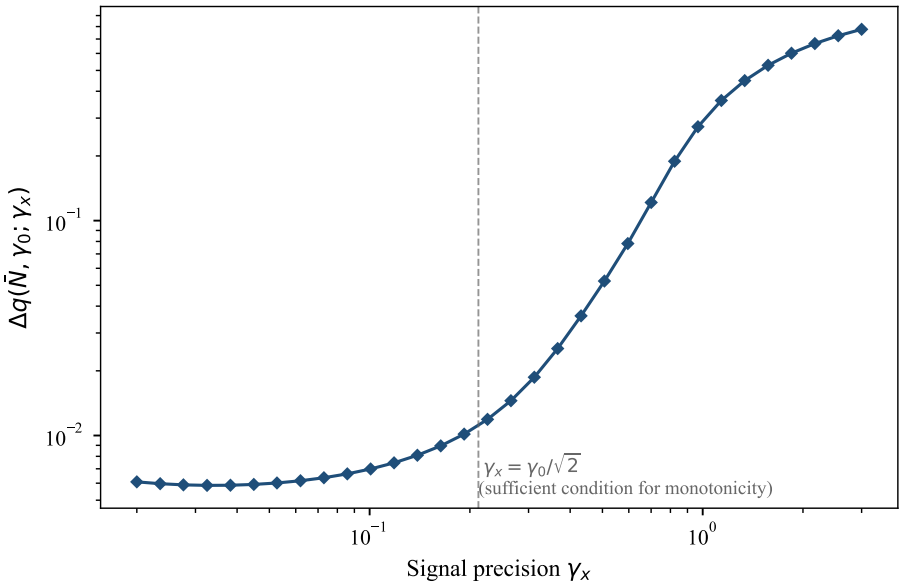
$$\frac{B(\gamma;\delta,a)}{B(\gamma+\delta;\delta,a)} = 1 + \frac{2\delta}{\gamma},$$

gives (A.14). Because

$$\log\left(1+\frac{2\delta}{\gamma}\right)$$

is strictly decreasing in γ, the proportional event effect is strictly larger at γ^L than at γ^H for all sufficiently large a. This proves (A.13). □

Figure 7: Numerical illustration of signal precision and the first resolution's effect

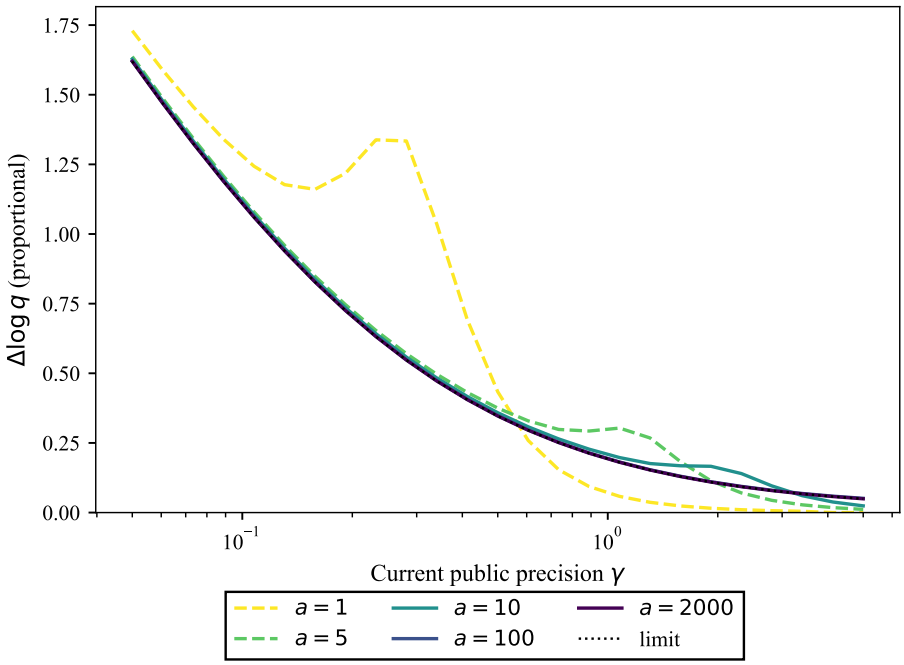


Notes. The curve shows $\Delta_0 q_{t+1}(\gamma_t, N_t; \delta)$ at $N_t = 100$, $\gamma_t = 0.3$, $a = 2$, and $G = U[0.1, 1]$. The dashed line marks $\delta = \gamma_t/\sqrt{2}$: any two signal precisions to its right satisfy the sufficient pairwise condition for the current-state result in Appendix A.9.1. The figure's γ_x denotes δ.

A.11 Additional Numerical Illustrations

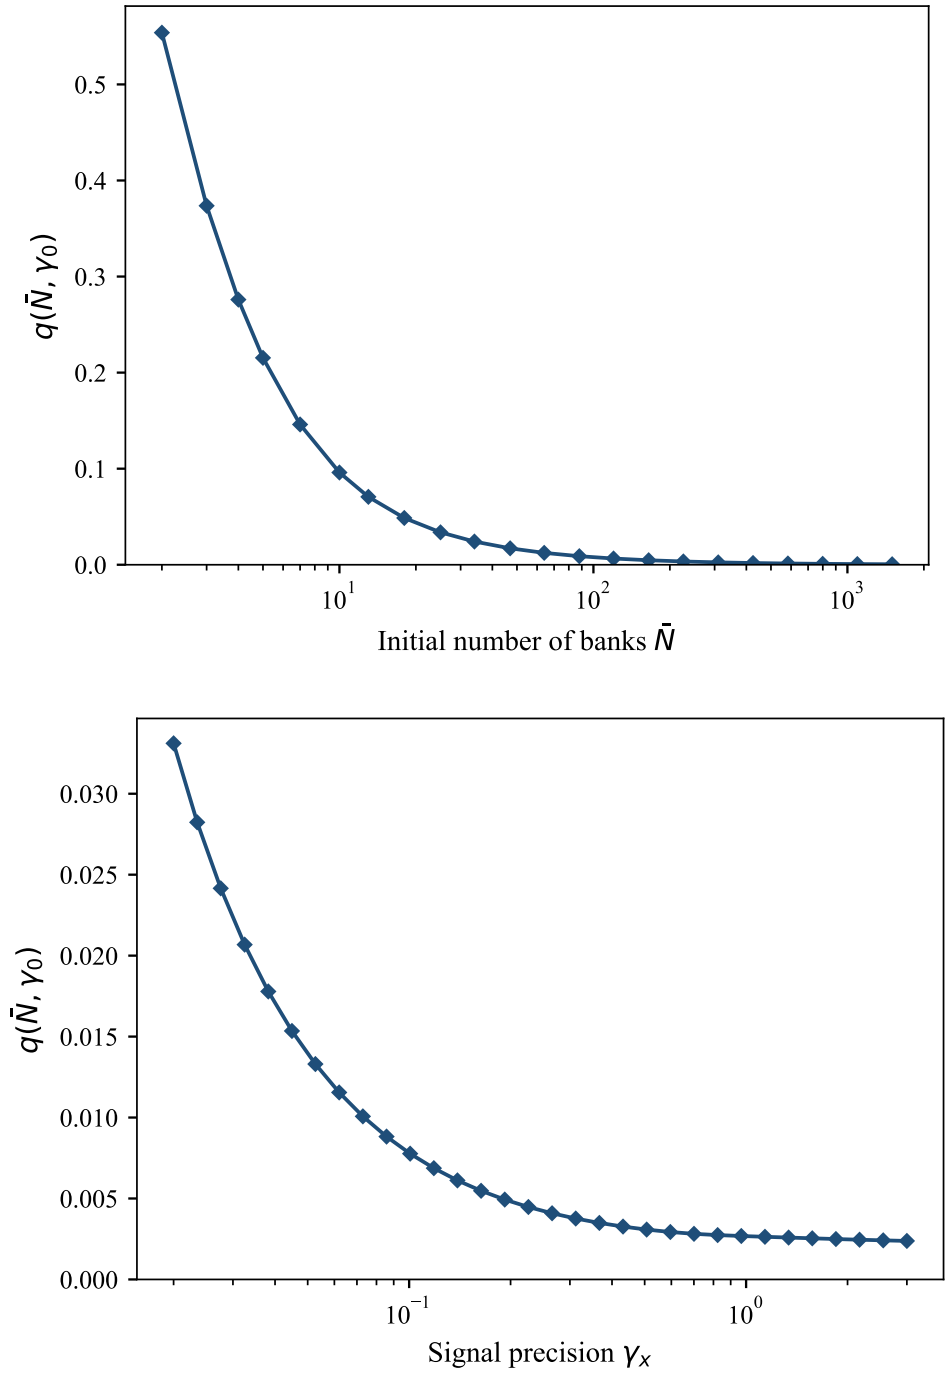
The following figures illustrate comparative statics of current resolution probabilities, expected resolution delays, and sensitivity to the waiting-cost distribution. The exercises use illustrative parameter values.

Figure 8: Initial public precision and the proportional event effect



Notes. The curves show $\Delta_0 \log q_{t+1}$ as initial precision varies, at $N_t = 100$, $\delta = 0.1$, and $G = U[0.1, 1]$, for the values of a in the legend. The dotted curve is the large-a limit $\log((N_t - 1)/(N_t - 2)) + \log(1 + 2\delta/\gamma_t)$. The decreasing profiles at large a illustrate the $n = 0$ comparison in Proposition 6; the small-a profiles show why its qualification matters.

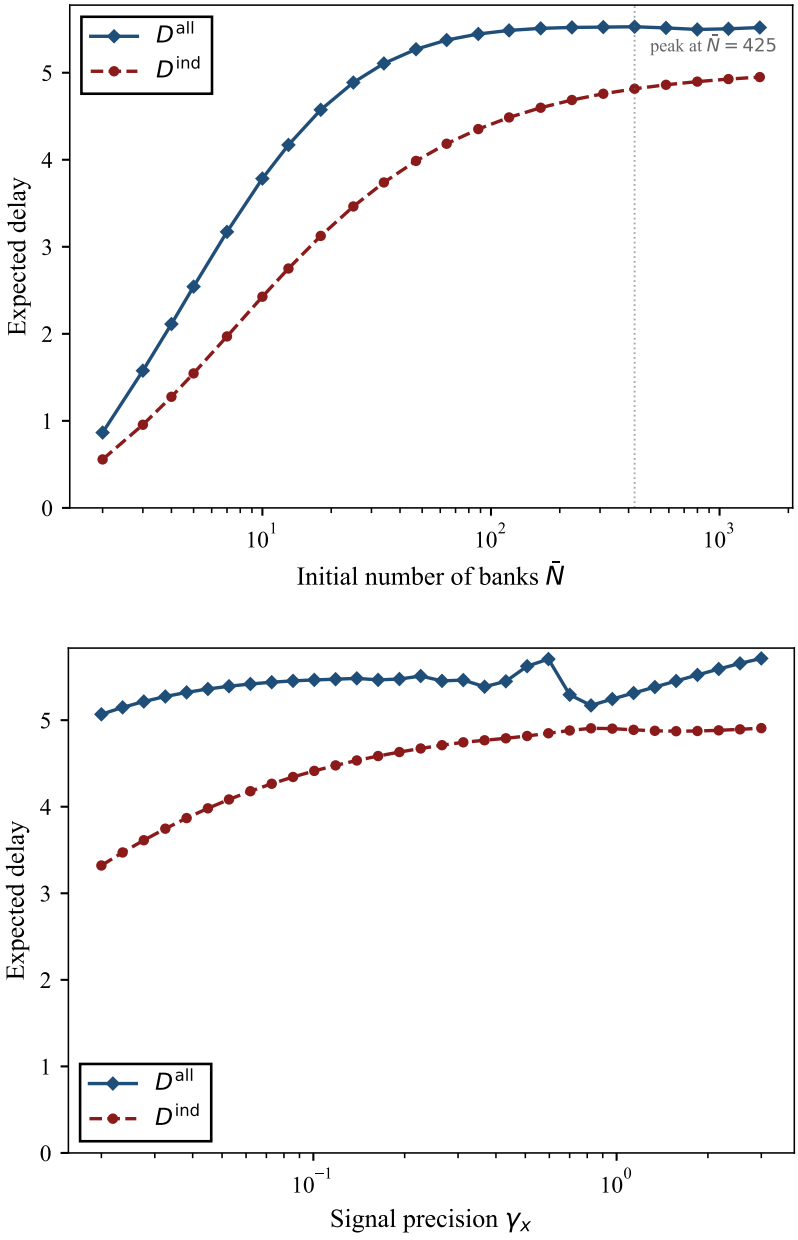
Figure 9: Numerical comparative statics of the current equilibrium resolution probability



Notes. The upper panel varies N; the lower panel varies δ. Unless varied on the horizontal axis, parameters are $N = 100$, $\gamma_0 = 0.3$, $\delta = 0.1$, $a = 2$, and $G = U[0.1, 1]$. These comparisons illustrate Parts 1 and 3 of Proposition 10. The figure labels $q(\bar{N}, \gamma_0)$ and γ_x correspond to $q^*(\gamma_0, \bar{N})$ and δ, respectively.

Figure 10: Market size, signal informativeness, and resolution delay

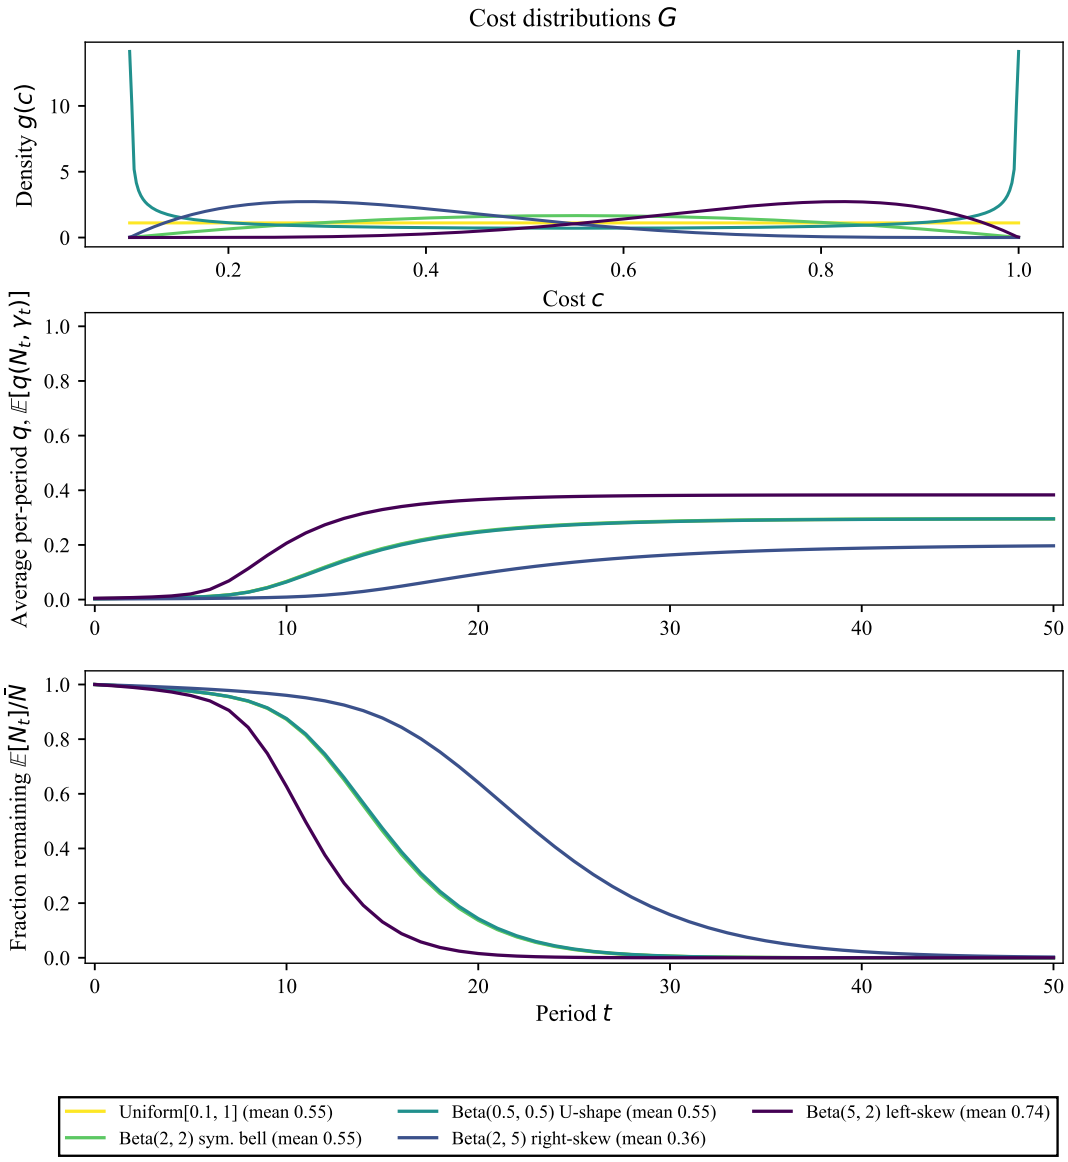
Figure 10 varies market size and signal precision and reports expected resolution delays. The time to resolve all loans is nonmonotone in both exercises, illustrating why the analytical orderings of current resolution probabilities do not imply corresponding orderings of total delay.



Notes. The upper panel varies initial market size; the lower panel varies signal precision. Unless varied on the horizontal axis, parameters are $\bar{N} = 100$, $\gamma_0 = 0.3$, $\delta = 0.1$, $a = 2$, and $G = U[0.1, 1]$. Both panels report D^{all} and D^{ind}, with immediate resolution assigned zero delay. These profiles illustrate the competing effects discussed after Proposition 10. The figure's γ_x denotes δ.

Figure 11: Sensitivity to the shape of the waiting-cost distribution

Figure 11 changes the shape of the waiting-cost distribution on a common support. Slow initial resolution followed by acceleration appears under each distribution, while the timing varies substantially, showing that the numerical duration of delay depends on cost heterogeneity.



Notes. All distributions have support $[0.1, 1]$, with $\bar{N} = 100$, $\gamma_0 = 0.1$, $\delta = 0.02$, and $a = 2$. Panels show cost densities, surviving-bank hazards $h_t = \mathbb{E}[n_t]/\mathbb{E}[N_t]$, and expected fractions remaining $\mathbb{E}[N_t]/\bar{N}$; the expected-q label denotes this survivor-weighted hazard. Paths are calculated from the equilibrium transition matrix. The uniform and two symmetric beta distributions share mean 0.55; the skewed distributions also change the mean, so those comparisons do not isolate shape at a fixed mean.

B Additional Data Documentation

This appendix documents the loan histories and empirical measures used in Section 4. It distinguishes recorded events from inferred boundaries of a loan's history and explains how both enter panel construction and outcome measurement.

B.1 Lender and Property Linkage

Lender-to-bank matching. The lender crosswalk combines names, company codes, lender location when available, and property location. A bank may appear under several abbreviations or legal-name variants, while unrelated banks may share names such as "First National Bank" or "Community Bank." The candidate list includes active and inactive institutions, branches, and subsidiaries from FDIC, NIC, and bank holding company reference records. I standardize names and abbreviations, classify lender entity types, and generate exact and approximate name matches.

Candidate assignments are checked against institutional operating dates, FDIC deposit coverage, county lending footprints, parent-bank CRE holdings, and related lender names appearing on the same property. When lender location is missing, property location provides weaker geographic evidence. Ambiguous names receive manual review, and the remaining candidates are resolved using match strength and consistency across property histories and regions. The crosswalk retains confidence information and links matched institutions to regulatory identifiers. Quarterly parent mappings account for changes in corporate ownership, including mergers, so parent attribution can change over the panel. This quarterly attribution is distinct from bank controls drawn from the fixed 2019Q4 regulatory snapshot.

Property classification and CMBS screening. CoreLogic's land-use and property-type codes can be broad or outdated. The category "commercial" can include multifamily properties, and some single-family properties also carry a commercial code. I refine these classifications using InfoGroup establishment listings, which report business locations and industries from 1990 onward. I geocode establishments and match them to properties, using business composition where it clearly identifies a property's use. For example, a property is classified as office when more than 70% of its businesses are office tenants.

I supplement these matches with property information from CompStak leases and Trepp CMBS records. For remaining unclassified properties with multiple records, I use large language models to verify property type individually. I classify the remaining properties using a machine-learning algorithm with CoreLogic's original codes, property characteristics, and local characteristics as inputs.

Residential and multifamily properties, agricultural land, vacant land, and exempt uses are outside this commercial-property universe.

The Trepp screen flags a mortgage when its origination falls within the observed coverage interval of a linked CMBS property. The economic loan inherits the flag from its originating record. The flag therefore depends on the mortgage's origination date as well as the property match. Loans carrying it are excluded from the bank-loan sample.

B.2 Loan Identity, Boundaries, and Quarterly Histories

Linking documents to an economic loan. I organize mortgage records chronologically within each property and remove duplicates. Modifications and refinancing records are linked to plausible prior mortgages using lender identity, dates, amounts, and lien information. Multiple concurrent loans on the same property retain separate histories. Linked modifications and same-lender refinancing preserve the economic loan identifier; refinancing by a different lender ends the prior relationship and creates a new identifier. When the original filing is missing, ownership transfers and later mortgage records can support an inferred origination, which remains flagged as inferred.

Contractual maturity and termination. Scheduled maturity uses the reported date where available. For missing dates, I use reported terms, then average terms by loan type and origination vintage, then loan-type averages, and finally a ten-year fallback. Later mortgage activity can establish that a loan continued beyond its imputed maturity. When a modification occurs after that date, I extend the imputed maturity to the nearest common loan term: 1, 2, 3, 5, 10, 15, 20, or 30 years.

Loan endings draw on default episodes, refinancing links, mortgage releases, successor loans, and qualifying sales. When stronger evidence of termination is unavailable, maturity-based boundaries and the observation window delimit the history. A loan that passes its estimated maturity without later activity becomes inactive in subsequent quarters. These administrative boundaries remain separate from the repayment and default events used as outcomes.

Quarterly expansion and event precedence. I expand each history into its active quarters within 2015Q1–2025Q2, including quarters without a new filing and the quarter of termination. When a successor replaces a loan, its quarterly history begins after the predecessor's final quarter to avoid overlapping exposure from the replacement. Amounts and rates can update in the event quarter; maturity and loan-status changes take effect in the following quarter. Same-lender refinancing can reset contractual age while preserving the economic loan relationship.

The panel assigns one primary event per loan-quarter. The regression event categories are derived from this primary event. A terminating event takes precedence over an extension recorded in the same quarter.

B.3 Event Classification and Outcome Aggregation

Extensions. A linked modification qualifies as an extension if its reported maturity exceeds the prior mortgage's maturity,. These rules identify continued financing when both contractual maturity dates are not observed from the implied maturity measures. Same-lender refinancing qualifies when the new maturity is later than the previous maturity. The extension outcome combines qualifying modifications and qualifying same-lender refinancing.

Figure 12 shows a recorded loan extension agreement.

Defaults. Default episodes combine ownership-transfer records, Property Basic foreclosure information, and episodes inferred from subsequent REO sales. I reconcile duplicates in that order

Figure 12: Example of a recorded loan extension agreement

1. **Acknowledgement of Balance Due**. Borrowers acknowledge and agree that the balance due on the Loan as of the Effective Date is as follows:

Principal	$9,694,731.36
Interest is paid in full	0.00
Balance Due	**$9,694,731.36**

2. **Maturity Date**. The Maturity Date of the Note is hereby modified and amended extended to January 5, 2011. Any reference in the Note or any of the other Loan Documents to the Maturity Date shall mean January 5, 2011.

Notes. The agreement acknowledges an outstanding principal balance of approximately $9.69 million and extends the note's maturity to January 5, 2011. It illustrates the recorded evidence used to identify continuation of an existing lending relationship.

of source priority: repossession, foreclosure filing, auction, and REO sale; the final category also includes short-sale evidence. The default event will terminate all the active loan associated with that property unless later record indicates continuation of some loans (e.g. refinance, modifications).

Sale payoffs and other repayments. Sale payoffs and other repayments follow the main-text definitions. Other repayments also include endings inferred from successor financing. They remain loan-ending events but are excluded from the informative-resolution outcome for Prediction 3. Table 7 summarizes the mapping.

From quarterly events to loan outcomes. The H3 resolution indicator records any default or sale payoff during 2022Q1–2025Q2 and is formed before the later observation filters used to select each loan's regression record. The extension indicator is the maximum qualifying extension indicator over the retained later-window observations. Both derive from the primary event assigned in each quarter. A loan with no qualifying event receives zero; the extension and resolution indicators need not sum to one.

B.4 Cohort, Peer Exposure, and Estimation Restrictions

Initial market stock comprises loans present in 2020Q1. For the later-window sample, I retain eligible observations from 2022Q1 onward and ultimately select one observation per economic loan.

Earlier informative-event counts sum defaults and sale payoffs during 2020Q1–2021Q4 within each county–property-type market. I calculate counts for the market and each parent bank, then subtract the loan's own parent bank's count from the market total. The initial peer stock is constructed in the same way from loans present in 2020Q1. Parent identities follow the quarterly bank mapping: the initial stock uses a fixed-quarter snapshot, while events use the parent identity recorded in their panel quarter. The peer resolution rate divides the informative-event count by the initial peer stock.

Regression preparation excludes CMBS-flagged observations and those without the required bank-capital classification. The preparation stage requires loan age, origination LTV, loan amount,

Table 7: Mapping quarterly events to empirical outcomes

Category	Included events	Informative resolution
Extension	Extension modification; internal refinancing with later maturity	No
Default	Classified default-related ending	Yes
Sale payoff	Classified payoff associated with a property sale	Yes
Other resolution	External refinancing; release; successor-loan repayment	No

Notes. The classification concerns the primary event in a loan-quarter. Extension uses event codes 1 and 6. Informative resolution defines the Prediction 3 outcome and uses codes 3 and 4. Other loan-ending codes 5, 8, and 11 remain in the panel but are excluded from this outcome. Maturity-based endings, censoring, non-extension modifications, and non-extension internal refinancing are not resolutions in this mapping. A loan can experience an extension in one quarter and a resolution in a subsequent quarter.

WFH exposure, transaction counts, bank concentration, and the maturity measure. After forming the extension outcome, I retain the first remaining later-window observation for each loan.

Each specification additionally requires an initial peer stock of at least ten, observed regression variables and fixed-effect identifiers, and available instruments for the joint IV results. Samples are held constant across columns within each table, with any further exclusions required by the PPML estimator.

B.5 WFH Exposure and Control Variables

WFH exposure. The inputs to equation (4.1) are matched at the three-digit North American Industry Classification System (NAICS) level. I use the industry intersection of County Business Patterns, CompStak property-use shares, and the WFH feasibility measure of Dingel and Neiman (2020). The employment and property-use weights are normalized to form a weighted average, and each county–property-type cell requires at least ten matched industries. Hotel properties use the other-commercial control inputs while retaining their separate property-type classification for loan events.

Age and maturity. Loan age uses the contractual origination date recorded in the panel and is measured in quarters at the first retained loan observation. It enters in natural logarithms. The maturity control is the minimum remaining quarters to maturity over the retained later-window history and enters without a logarithmic transformation, including eligible zero and negative values. Contractual updates follow the timing in Appendix B.2.

Origination LTV. Origination LTV divides the mortgage amount assigned to the loan by a constructed collateral value. The value draws on recorded sales and assessor values. I screen unusable prices, adjust loan amounts that cover multiple properties, and use property price indices to bring earlier values forward to origination. Assessor candidates include reported market values and assessed values adjusted for fractional assessment. Among candidates yielding an admissible LTV,

I favor the larger assessor value, with calculated total value as a fallback. The construction can use observations after origination. The regressions use LTV anchored to acceptable observed values; separate loan-based collateral-value imputation does not fill LTV when such an anchor is missing.

Liquidity, concentration, and market size. Pre-pandemic transaction volume is the average annual number of qualifying property purchases during 2015–2019. The transaction-count control is the natural logarithm of one plus this average. Annual turnover divides the average by the number of properties in the relevant CoreLogic commercial-property universe, including properties without active mortgages. Bank concentration is the Herfindahl index of parent banks' shares of market loan amounts at 2019Q4. All specifications control for initial peer stock; the full-control columns of Predictions 1–3 also include the informative-resolution rate.

B.6 Disclosure and Relative Transaction-Price News

The disclosure indicator assigns Idaho, Kansas, Louisiana, Mississippi, Montana, New Mexico, Texas, Utah, Alaska and Wyoming to the nondisclosure group. All other states enter the disclosure group, including Missouri, which is separately identified as partial disclosure. Transaction-price availability varies within the disclosure group.

For sale-payoff event e, let P_e denote its transaction price and A_e its assessed-value benchmark, adjusted for fractional assessment and brought forward to the sale date using regional and property-type price indices. I anchor this adjustment at the most recent transaction at which the assessed value was aligned with the transaction price. Both values are in dollars. The price gap is

$$g_e = \log P_e - \log(1 + A_e). \tag{B.1}$$

I classify news relative to the pooled mean gap among sale-payoff events with observed prices in 2020Q1–2021Q4. Events above that mean are good news; events at or below it are bad news. Events missing a price or valuation input enter neither group.

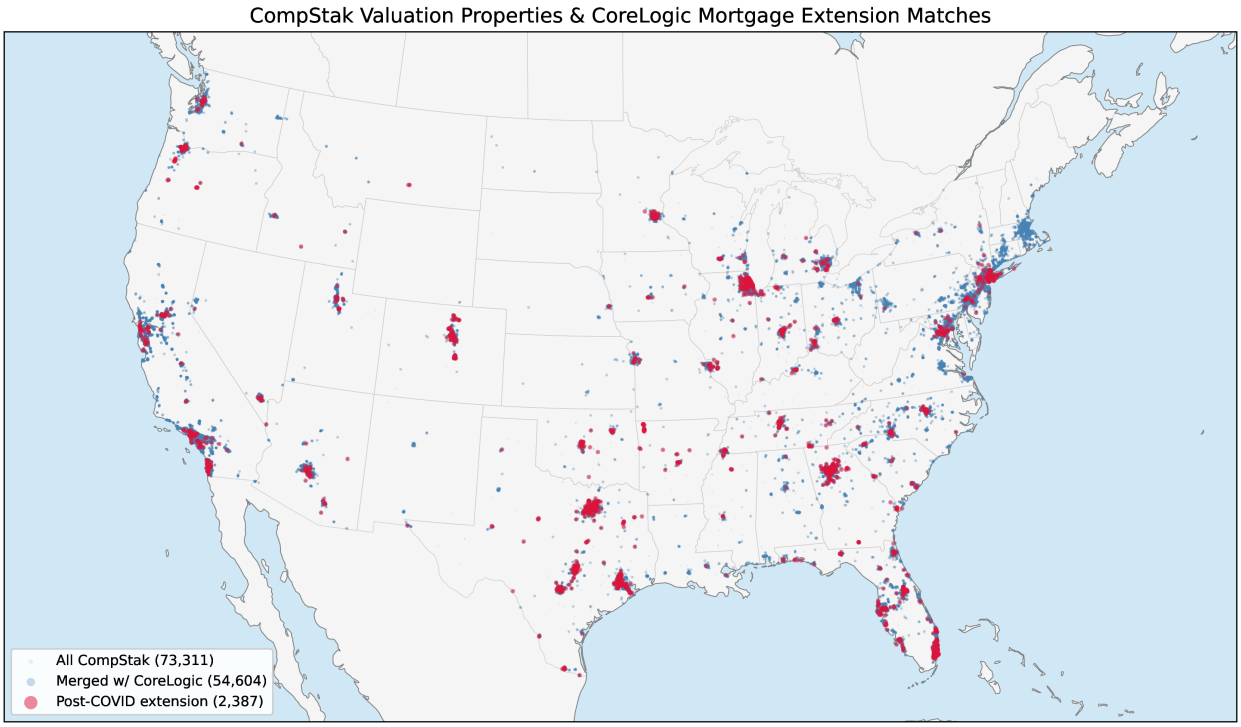
I aggregate good- and bad-news counts within county–property-type markets, excluding the loan's own parent bank. The comparison uses loans in disclosure states and a common scaling factor for the two counts. The pooled cutoff is calculated before restricting the final regression to disclosure states. The measure therefore captures relative transaction-price news against an assessed-value benchmark.

B.7 Coverage Comparison

Panel A of Figure 6 compares parent-bank origination totals for the 626,839 sample loans with regulatory nonfarm nonresidential real estate loan balances in 2020Q1 (F160 + F161). Regulatory amounts reported in thousands of dollars are converted to dollars before taking logarithms. Sampled loan amounts remain at origination values, without subsequent amortization adjustments.

Of the 2,900 represented parents, 2,783 have positive matched amounts on both axes and enter Panel A of Figure 6. The remaining 117 comprise 107 parents without a matched regulatory record,

Figure 13: CompStak valuation properties and CoreLogic matches



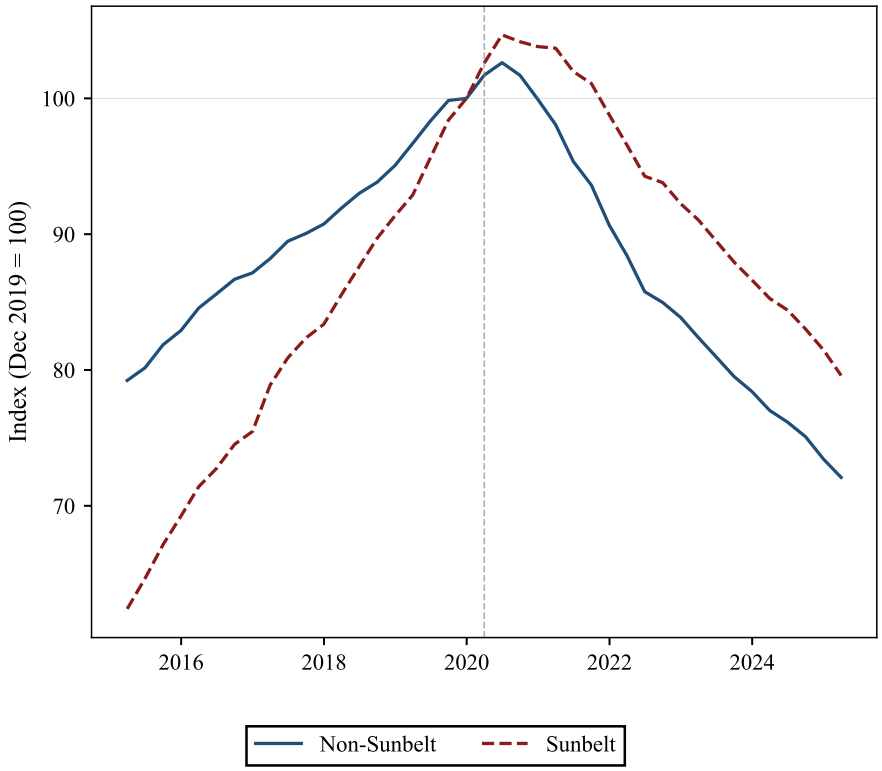
Notes. Counts refer to properties in the valuation exercise. Table 1 reports the separate loan-level regression sample.

one with a missing CRE balance, and nine with zero CRE balances. The regulatory measure uses reported CRE loan balances without adding a separate CMBS amount.

B.8 CompStak Coverage

Figure 13 shows the CompStak valuation-property dataset and its overlap with CoreLogic mortgage records. Of 73,311 CompStak properties, 54,604 are matched to CoreLogic; 2,387 of these matched properties are associated with a post-COVID mortgage extension.

Figure 14: Office Leasing Revenue



Notes. Total annualized revenue is the sum of leased square footage times annual effective rent for leases active at each observation date. Revenue is expressed in December 2022 dollars and indexed to December 2019 = 100. The figure displays 2015Q1–2025Q1; the vertical line marks 2020Q1. The Sunbelt classification follows the existing CompStak market grouping. Source: CompStak.

C Additional Background and Motivating Evidence

C.1 Office Leasing Revenue

Figure 14 plots total real leasing revenue for office properties in Sunbelt and non-Sunbelt markets in the CompStak sample.

C.2 Descriptive Extension Shares

The extension figures summarize annual recorded events for bank-held office loans. The mortgage-record sample identifies office properties using the InfoGroup classification, excludes CMBS-linked collateral, and retains records matched to banks with an observed 2019Q4 capital buffer. Lenders are assigned to their fixed 2019 parent banks. This descriptive sample differs from the economic-loan cohort in Section 5, where the panel retains loans without an event and parent attribution can change over time.

The numerator counts modification records tagged as clear extensions. The denominator adds refinancing records and matched default-event counts in the same year and bank group. Refinancing

includes external, internal, and unclassified refinancing; its internal-maturity subcategories are not counted again. Sale payoffs are excluded from this denominator. An internal refinancing that extends maturity therefore remains a refinancing in these figures, while it contributes to the extension outcome in the loan-level regressions. The plotted shares are ratios of pooled event counts, not averages of bank-specific shares or probabilities for the stock of outstanding loans. The year 2025 is incomplete.

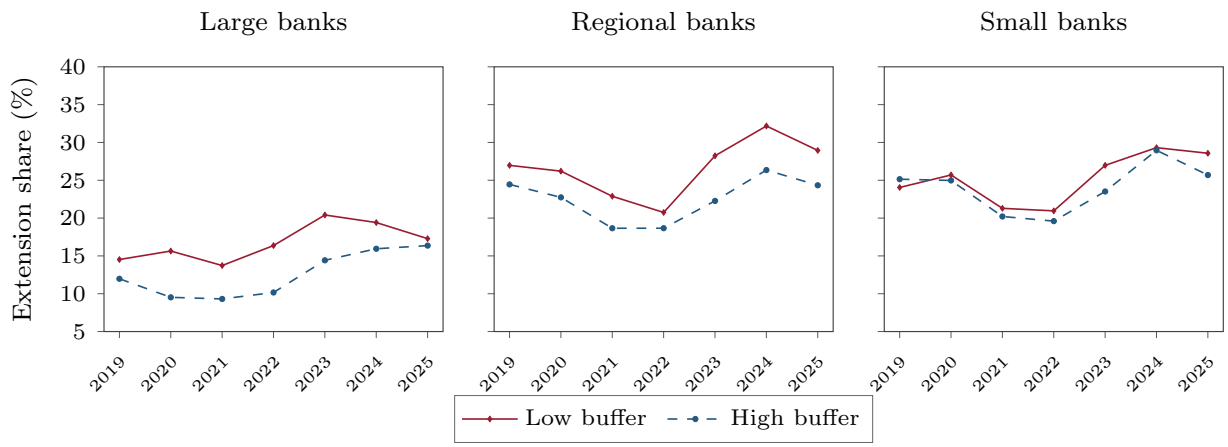
The capital buffer equals the mark-to-market CET1 ratio minus the applicable CET1 requirement, expressed in percentage points. Both the fixed-cutoff split and the median comparison use its 2019Q4 value. Median cutoffs are calculated separately by bank size among CRE-active mortgage-parent banks with positive buffers.

Figure 2 separates banks at asset cutoffs of $100 billion and $1.284 billion and divides each size group at a five-percentage-point capital buffer. Figures 15 and 16 provide the median-buffer and large-bank comparisons.

C.3 Alternative Capital-Buffer Split

Figure 15 shows that extension shares rise after 2022 on both sides of the within-group median capital buffer. The broad pattern therefore does not depend on placing every group at a common five-point cutoff.

Figure 15: Extension Shares Using Within-Group Median Capital Buffers

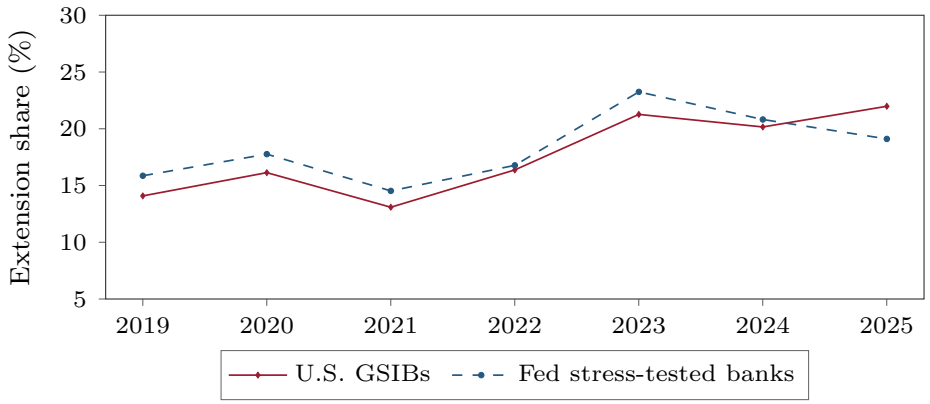


Notes. Low-buffer banks are at or below the size-specific median; high-buffer banks are above it. The medians are 4.29 percentage points for large banks, 6.37 for regional banks, and 8.65 for small banks. Sample, median-construction, and extension-share definitions follow Appendix C.2. In this supporting export, the large-bank group includes banks whose maximum observed assets reach $100 billion; among the remaining banks, the regional/small split uses average assets above/at or below $1.564 billion. The displayed window is 2019–2025. Sources: CoreLogic and FFIEC.

C.4 GSIBs and Banks in the Stress-Tested Sample

Figure 16 shows that extension shares also rise after 2022 among GSIBs and banks in the stress-tested sample.

Figure 16: Extension Shares among GSIBs and Stress-Tested Banks



Notes. GSIB membership is fixed at 2019Q4. Of the 32 banks in the Federal Reserve's 2024 capital-requirements list, 30 are matched to regulatory identifiers and 27 appear in the mortgage sample, including all eight U.S. GSIBs. These are counts of distinct banks across the sample, rather than annual counts. Both series use the sample and extension-share definitions in Appendix C.2. The figure displays 2019–2025. Sources: CoreLogic, FFIEC, and Federal Reserve (2024).

Table 8: Sensitivity to the Window Cutoff

Quarter Cutoff	2021Q2 (1)	2021Q3 (2)	2021Q4 (3)	2022Q1 (4)	2022Q2 (5)
Coefficient on earlier resolutions at other banks					
Baseline OLS	−0.268	−0.241	−0.277**	−0.246**	−0.250***
	(0.200)	(0.154)	(0.136)	(0.111)	(0.096)
Adding bank FE	−0.265**	−0.248***	−0.267***	−0.232***	−0.241***
	(0.112)	(0.074)	(0.078)	(0.064)	(0.058)
Adding local controls and R	−0.234**	−0.228***	−0.256***	−0.229***	−0.229***
	(0.116)	(0.075)	(0.078)	(0.065)	(0.055)
Mean extension rate (%)	3.79	3.56	3.31	3.05	2.78
Observations	674,188	651,724	626,839	606,900	585,321

Notes. The cutoff is the last quarter of the earlier window. The outcome is extension during the subsequent window, ending in 2025Q2. Coefficients and standard errors are in percentage points; counts use the fixed total-event SD of 112.9793. Samples are specific to each cutoff. Each specification includes county and property-type fixed effects, loan characteristics, WFH exposure, and initial peer-stock controls. The second row adds parent-bank fixed effects; the third also adds local controls and the peer resolution rate. The 2021Q4 column uses the reconstructed outcome window for this exercise. Standard errors clustered by county are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

D Additional Empirical Results

D.1 Sensitivity to the Window Cutoff

Table 8 examines whether the negative relationship in Section 5.1 depends on the boundary between the earlier and later windows. The earlier window starts in 2020Q1 and ends at each reported cutoff; the later window begins in the following quarter and ends in 2025Q2. Each column uses loans that survive into its later window. Outcomes and earlier peer exposures are reconstructed for each boundary, with 2021Q4 serving as this exercise's reconstructed benchmark.

The full-control count coefficient remains negative and statistically significant at the 5 percent level across all five cutoffs, ranging from −0.256 to −0.228 percentage points per approximately 113 earlier peer resolutions. The samples and extension frequencies change with the length and starting date of the later window, while the count scale is held fixed.

D.2 Joint and Default-Only Disclosure Comparisons

This subsection supplements the sale-payoff comparison in Section 5.2. I first include sale payoffs and defaults jointly to examine whether the disclosure differential is specific to sales. I then report a default-only specification.

The joint specification includes both types of earlier informative resolutions and their disclosure

Table 9: Disclosure in the Joint Sale–Default Specification

	(1) Baseline	(2) + Bank FE	(3) + Local controls and rate
Sale payoffs × disclosure	−1.391**	−0.812**	−0.905**
	(0.601)	(0.385)	(0.407)
Defaults × disclosure	5.402*	2.215	2.513
	(3.273)	(1.536)	(1.684)
Sale payoffs (SD)	0.330	−0.212	−0.092
	(0.554)	(0.319)	(0.337)
Defaults (SD)	−5.293	−2.065	−2.331
	(3.258)	(1.512)	(1.655)
WFH (SD)	0.307**	0.355***	0.411***
	(0.148)	(0.136)	(0.144)
Resolution rate R (pp)			−0.016
			(0.012)
Parent-bank fixed effects	No	Yes	Yes
Local controls	No	No	Yes
Initial peer-stock controls	Yes	Yes	Yes
Peer resolution rate	No	No	Yes
Observations	626,839	626,839	626,839
R^2	0.0451	0.0728	0.0729

Notes. Outcome: extension during 2022Q1–2025Q2. Coefficients and standard errors are in percentage points. Earlier peer events occur during 2020Q1–2021Q4 and exclude the loan's own parent bank. All count regressors use the common total-event SD of 112.9793; rates are in percentage points. Sale and default counts and their disclosure interactions enter jointly. The disclosure main effect is absorbed by county fixed effects. All columns include county and property-type fixed effects, log loan age, origination LTV, raw minimum quarters to maturity, and standardized WFH exposure. Local controls are pre-pandemic turnover, log transaction volume, and bank concentration. All columns include initial peer-stock controls; column 3 also includes the peer resolution rate. Standard errors clustered by county are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

interactions:

$$\begin{aligned} Ext_{i,c,p,b,t_2} = \alpha &+ \beta_S SaleCounts_{c,p,-b,t_1} + \beta_D DefaultCounts_{c,p,-b,t_1} \\ &+ \theta_S (SaleCounts_{c,p,-b,t_1} \times D_c) \\ &+ \theta_D (DefaultCounts_{c,p,-b,t_1} \times D_c) \\ &+ \Theta' W_{i,c,p,b,t_0} + \lambda_c + \kappa_p + \psi_b + \varepsilon_{i,c,p,b,t_2}. \end{aligned} \tag{D.1}$$

The controls follow the progression in the main sale-payoff specification: initial peer-stock controls in every column, parent-bank fixed effects from column 2, and local controls and the peer resolution rate in column 3. Both event types are scaled by the common total-event standard deviation.

Table 9 reports a negative sale-payoff–disclosure interaction in every column. In the full specification, the sale interaction is -0.905 ($p = 0.026$), while the default interaction is 2.513 ($p = 0.136$). Their difference, the sale-minus-default interaction, is -3.418 percentage points ($p = 0.094$). This contrast measures $\theta_S - \theta_D$ on the common count scale.

Table 10: Disclosure Slopes in the Joint Sale–Default Specification

	(1) Baseline	(2) + Bank FE	(3) + Local controls and rate
Sale payoffs			
Nondisclosure slope	0.330	−0.212	−0.092
	(0.554)	(0.319)	(0.337)
Disclosure slope	−1.061***	−1.024***	−0.997***
	(0.269)	(0.251)	(0.266)
Disclosure differential	−1.391**	−0.812**	−0.905**
	(0.601)	(0.385)	(0.407)
Defaults			
Nondisclosure slope	−5.293	−2.065	−2.331
	(3.258)	(1.512)	(1.655)
Disclosure slope	0.109	0.150	0.182
	(0.218)	(0.199)	(0.203)
Disclosure differential	5.402*	2.215	2.513
	(3.273)	(1.536)	(1.684)
Interaction difference	−6.793	−3.026	−3.418
	(3.831)	(1.876)	(2.043)
Equality-test p-value	0.0764	0.1068	0.0944

Notes. Linear combinations from Table 9. Slopes and differentials are percentage points per 112.9793 events, the common total-event SD. Standard errors for the disclosure-state slopes use the estimated covariance between each main effect and interaction. The interaction difference is the sale minus default disclosure interaction on the common count scale. Standard errors clustered by county are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 10 reports slopes within each disclosure regime. The disclosure-state slopes are $\beta_S + \theta_S$ for sales and $\beta_D + \theta_D$ for defaults; their standard errors use the estimated covariance between each main effect and its interaction. In column 3, the slopes are -0.997 ($p < 0.001$) for sales and 0.182 ($p = 0.372$) for defaults.

The default-only specification omits sale-payoff counts and their disclosure interaction from equation (D.1). Table 11 reports a full-control disclosure interaction of 1.179 ($p = 0.094$). The default slope within disclosure states is -0.388 (standard error 0.142, $p = 0.006$). The main information-content result remains the consistently negative sale interaction.

Table 11: Disclosure in the Default-Only Specification

	(1) Baseline	(2) + Bank FE	(3) + Local controls and rate
Defaults × disclosure	2.767**	1.252**	1.179*
	(1.407)	(0.636)	(0.703)
Defaults (SD)	−3.311**	−1.693**	−1.568**
	(1.478)	(0.681)	(0.756)
WFH (SD)	0.205	0.260**	0.354**
	(0.143)	(0.132)	(0.140)
Resolution rate R (pp)			−0.027**
			(0.012)
Parent-bank fixed effects	No	Yes	Yes
Local controls	No	No	Yes
Initial peer-stock controls	Yes	Yes	Yes
Peer resolution rate	No	No	Yes
Observations	626,839	626,839	626,839
R^2	0.0451	0.0728	0.0728

Notes. Outcome: extension during 2022Q1–2025Q2. Coefficients and standard errors are in percentage points. Earlier peer events occur during 2020Q1–2021Q4 and exclude the loan's own parent bank. All count regressors use the common total-event SD of 112.9793; rates are in percentage points. The disclosure interaction compares the default-count slope across disclosure regimes. The disclosure main effect is absorbed by county fixed effects. All columns include county and property-type fixed effects, log loan age, origination LTV, raw minimum quarters to maturity, and standardized WFH exposure. Local controls are pre-pandemic turnover, log transaction volume, and bank concentration. All columns include initial peer-stock controls; column 3 also includes the peer resolution rate. Standard errors clustered by county are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

D.3 Initial Uncertainty and Extension in Levels

The main test in Section 5.3 concerns proportional changes in resolution probability. This supporting regression examines whether WFH exposure strengthens the relationship between earlier resolutions and extension probability in levels:

$$\begin{aligned} Ext_{i,c,p,b,t_2} = \alpha &+ \beta_H (ResCounts_{c,p,-b,t_1} \times WFH_{c,p}) \\ &+ \beta_X ResCounts_{c,p,-b,t_1} + \rho WFH_{c,p} \\ &+ \Theta' W_{i,c,p,b,t_0} + \lambda_c + \kappa_p + \psi_b + \varepsilon_{i,c,p,b,t_2}. \end{aligned} \tag{D.2}$$

The scaled count is centered at its original reference mean, and WFH exposure is standardized. All columns include initial peer-stock controls; column 3 also includes local controls and the peer resolution rate.

Table 12 shows a negative interaction in every column.

Table 12: Initial Uncertainty and Extension Responses in Levels

	(1) Baseline	(2) + Bank FE	(3) + Local controls and rate
Resolution counts × WFH	−0.190***	−0.099*	−0.099*
	(0.063)	(0.051)	(0.054)
Resolution counts (centered)	−0.514***	−0.395***	−0.389***
	(0.150)	(0.112)	(0.119)
WFH exposure	0.343**	0.358***	0.394***
	(0.145)	(0.138)	(0.141)
Resolution rate R (pp)			−0.022*
			(0.012)
Parent-bank fixed effects	No	Yes	Yes
Local controls	No	No	Yes
Initial peer-stock controls	Yes	Yes	Yes
Peer resolution rate	No	No	Yes
Observations	626,839	626,839	626,839
R^2	0.0452	0.0729	0.0729

Notes. Outcome: extension during 2022Q1–2025Q2. Coefficients and standard errors are in percentage points. Earlier peer events occur during 2020Q1–2021Q4 and exclude the loan's own parent bank. All count regressors use the common total-event SD of 112.9793; rates are in percentage points. All columns include county and property-type fixed effects, log loan age, origination LTV, raw minimum quarters to maturity, and standardized WFH exposure. Local controls are pre-pandemic turnover, log transaction volume, and bank concentration. Counts are centered at the original reference mean before interacting with WFH. All columns include initial peer-stock controls; column 3 adds the peer rate. Standard errors clustered by county are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

D.4 Price-News Coverage and Scaling

The price-news test in Section 5.4 uses the transaction classification and peer-count construction in Appendix B.6.

All three columns use the same 544,506 loans, retained from 544,701 eligible observations after the common-sample restrictions, including singleton removal. Columns 1–2 count priced and unpriced sale payoffs. Column 3 divides price-observed sale payoffs into good- and bad-news events, which together exhaust the priced event set.

All-sale, good-news, and bad-news counts are divided by the same standard deviation of all peer sale-payoff counts in the final news sample, $s_{\text{news}} = 78.4594$. This denominator differs from the common total-event standard deviation of 112.9793 used in the H2 disclosure comparisons. WFH exposure retains its reference standardization from before the disclosure-state and news-specific sample restrictions.

In column 2, mean news is the average price gap among priced peer sales. I center and standardize it among sample loans with at least one such sale. Loans without priced peer sales remain in the sample and receive a standardized mean-news value of zero. Their all-sale count may still be positive because it includes unpriced events. Column 3 includes the good- and bad-news counts without a separate unpriced-event count.

The good-minus-bad contrast tests equality on the common sale-count scale: it is 0.020 percentage points (standard error 0.266, $p = 0.939$).