Selling Distress
The Economics of Failed Bank Resolutions

Andrew Perry*

DRAFT: April 15, 2026
DO NOT DISTRIBUTE

ABSTRACT

This paper introduces a novel dataset of bank resolution terms constructed from over forty years of FDIC press releases to analyze the determinants of franchise value in failed bank auctions. Using a structural model of the auction process that encompasses both the liability and asset sides of the failed bank's balance sheet, I characterize the deposit franchise as a wasting asset and endogenize the acquirer's optimal asset take-up decision as a function of comparative advantage relative to the FDIC's own resolution technology. The "Toxic Ratio", the proportion of assets retained by the FDIC in receivership, emerges as the dominant driver of resolution outcomes across the pricing, cost, and selection channels. Because the retained asset share and the deposit premium are jointly determined at the auction, I instrument the realized toxic ratio with the pre-failure loan-loss provision (LLP) ratio. The headline cost result is that a one percentage point increase in the loan-loss provision ratio raises the cost to the Deposit Insurance Fund by approximately 86 basis points of total assets, making asset quality the dominant quantitative determinant of fiscal exposure. Accelerated deposit runoff prior to failure further reduces bid premiums and increases resolution costs, consistent with the franchise-value mechanism. This trade-off highlights a shift in regulatory technology from liquidation to financial engineering, with significant implications for moral hazard and the "Too-Big-To-Fail" doctrine.

JEL classification: G21, G28, H12.

*Andrew Perry is with the Haas School of Business, U.C. Berkeley. Acknowledgments: I want to thank Darrith Phan and Kavita Selva for their amazing help collecting the public bank failure data that is novel to this project.

I. Introduction

Bank failures are rare, consequential, and poorly understood. When a bank fails, the Federal Deposit Insurance Corporation (FDIC) must rapidly resolve the institution, allocating losses among depositors, the insurance fund, and private acquirers. While the aggregate fiscal costs of these failures are well documented, the specific transaction mechanisms have largely remained opaque to researchers. This paper introduces a novel dataset of FDIC resolution terms constructed from over forty years of administrative press releases, enabling the first systematic analysis of how resolution mechanisms vary across failure types, economic conditions, and acquiring bank characteristics. The data reveal that resolution structure is far from mechanical: acquirer pricing, loss-sharing arrangements, and deposit franchise valuation reflect deliberate economic choices whose determinants this paper identifies and quantifies.

The stakes of efficient resolution are high. During the Savings and Loan (S&L) crisis of the 1980s, the collapse of the thrift industry generated resolution costs in excess of 150 billion dollars. More recently, the Global Financial Crisis and the regional banking failures of 2023 demonstrated that even a relatively small number of failures can impose significant stress on the Deposit Insurance Fund (DIF). Despite the magnitude of these events, the existing academic literature faces a significant data limitation. Most research relies either on aggregate cost estimates or on privileged internal data that is unavailable to the broader scientific community. Standard public regulatory filings such as Call Reports provide a detailed view of the bank prior to failure, but the data trail typically ends the moment the charter is revoked. This leaves researchers unable to observe the specific premiums paid by acquirers, the precise volume of assets retained by the receiver, or the structure of loss-sharing agreements that define the modern resolution.

This paper bridges this gap by introducing a hand-collected dataset constructed from the universe of FDIC press releases spanning the period from 1983 to 2024. By systematically encoding over forty years of resolution announcements, I transform qualitative legal disclosures into a structured economic dataset containing 1,860 observations. This data collection effort quantifies critical resolution terms that are absent from standard regulatory reports. These variables include the premia paid for deposit franchises, the magnitude of assets retained by the FDIC in receivership, and the precise structure of loss-sharing agreements. This granular view allows for the decomposition of resolution costs and provides a new lens through which to evaluate the determinants of franchise value in failed bank auctions.

The institutional framework governing these resolutions has evolved significantly over the sample period. The passage of the Federal Deposit Insurance Corporation Improvement Act of 1991 (FDICIA) marked a pivotal shift in the regulatory landscape. Designed to correct the incentive problems and regulatory forbearance that exacerbated the S&L crisis, FDICIA mandated a least-cost resolution test. This requirement forces the FDIC to select the resolution method that imposes the minimum expense on the insurance fund barring a systemic risk exception. Consequently, the resolution mix shifted dramatically from a landscape where forbearance was common to one dominated by Purchase and Assumption (P&A) transactions. In a P&A transaction, a healthy acquiring

institution purchases some or all of the failed bank's assets and assumes its deposit liabilities. This method is generally preferred as it preserves the franchise value of the failed institution and maintains the continuity of banking services for the local community.

The historical arc of these failures is visualized in Figure 1 which plots the frequency and resolution method of every bank failure in the sample. The data reveals three distinct epochs of distress: the S&L crisis of the late 1980s, the GFC of 2008, and the concentrated regional banking shock of 2023. While the sheer volume of failures in the earlier period dwarfs modern episodes, the visual composition of the bars highlights a profound shift in regulatory technology. The 1980s exhibit a fragmented approach where the FDIC frequently relied on direct Deposit Payoffs and Insured Deposit Transfers to close small, chaotic institutions. By contrast, the modern era is characterized by the monolithic dominance of the P&A transaction.

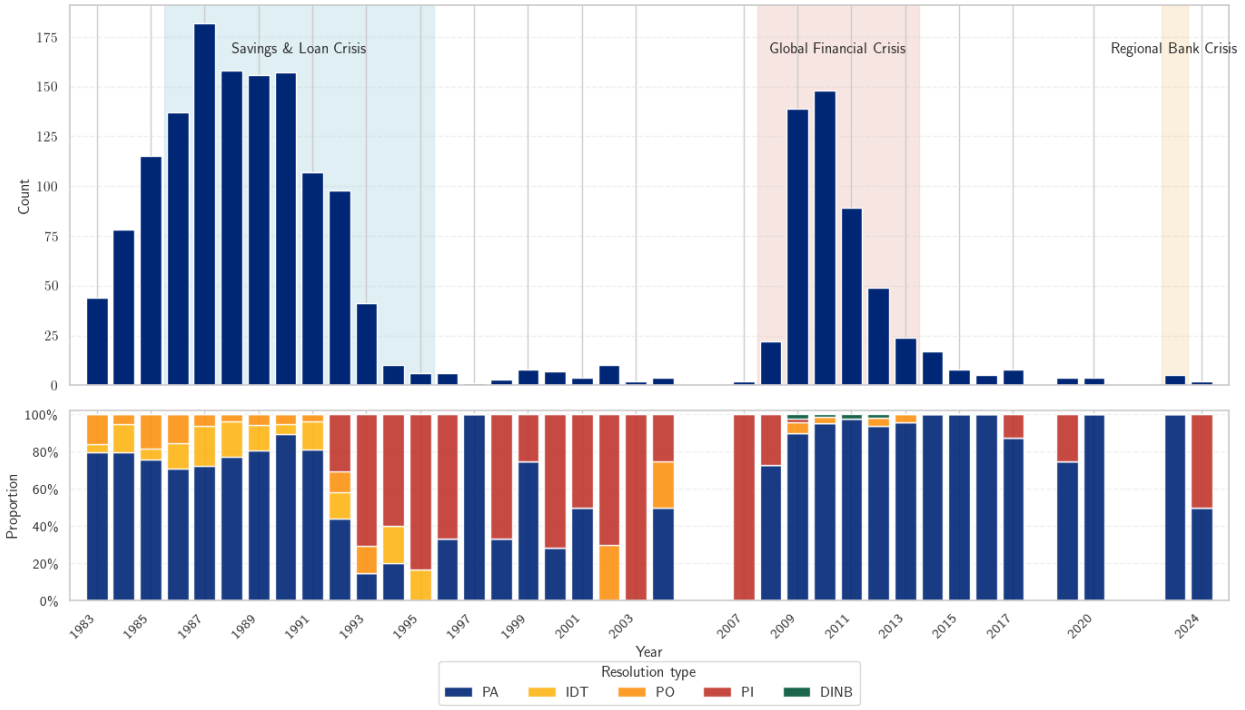


Figure 1. Annual bank failures and resolution mix, 1983–2024. The top panel shows the count of failures with shaded bands for the Savings & Loan crisis (blue), Global Financial Crisis (pink), and 2023 regional bank crisis (tan). The bottom panel gives the share of each FDIC resolution type (PA, IDT, PO, PI, DINB) by year.

This uniformity in the later years is deceptive. While Figure 1 suggests that the resolution process has become standardized, the underlying transaction structures have actually become more complex. The disappearance of the deposit payoff is not merely a function of bidder interest but a result of financial engineering. As the empirical results of this paper demonstrate, the ubiquity of the P&A transaction in the post-2008 period is largely sustained by the FDIC's increasing willingness to retain toxic assets. In the 1980s, a bank with a poor loan portfolio might have been liquidated. Today, that same institution is sold via a P&A transaction where the acquirer purchases the franchise but leaves the distressed assets behind with the receiver. Thus, the clean blue bars

of the modern era mask a significant transfer of risk that can only be observed by analyzing the specific asset retention terms.

To interpret the variation in the pricing of these transactions, I develop a structural model of the FDIC auction process that includes both the asset and liability side of the failed bank's balance sheet. Each bidder's marginal value for the failed bank's assets declines in the share acquired, reflecting portfolio heterogeneity, regulatory capital constraints, and the diminishing returns of operational absorption. Bidders optimally select an asset take-up share by equating their marginal asset valuation to the FDIC's own expected recovery value, partitioning the portfolio into a "transferable" tranche managed more efficiently by the private sector and a "retained" tranche that the FDIC works out in receivership. On the liability side, the deposit franchise is a wasting asset where value is driven by the stability and interest-rate sensitivity of the deposit base. The FDIC's statutory least-cost mandate induces a linear scoring rule that compresses the three-dimensional bid—deposit premium, asset purchase price, and asset take-up share—into a single scalar, preserving the full tractability of first-price auction theory. The model's central prediction is that the retained asset share and the deposit premium are positively related through two reinforcing channels. A mechanical composition effect, whereby a lower take-up share shifts the bid toward the premium line, and an equilibrium private-value effect, whereby shedding negative-surplus assets raises the winning bidder's willingness to pay for the cleansed franchise.

The empirical analysis utilizes the new dataset to test these theoretical predictions through three distinct channels. The first is the deposit pricing channel, which examines the bid premium paid for the deposit franchise conditional on a successful sale. The second is the FDIC cost channel, which analyzes the fiscal expense absorbed by the Deposit Insurance Fund. The third is the P&A selection channel, which models the probability that a bank is resolved via purchase and assumption rather than a direct deposit payoff.

Consistent with the theoretical framework, I find evidence that the stability of the deposit base matters. Banks that experience accelerated deposit runoff in the quarter prior to failure impose significantly higher costs on the FDIC. This finding supports the view that the destruction of franchise value prior to closure directly translates into fiscal losses. Furthermore, the analysis of deposit betas suggests a negative relationship between interest rate sensitivity and bid premiums although this result is statistically noisier than the runoff effect. This implies that while acquirers theoretically value low-beta deposits, the chaotic environment of a failed bank auction may obscure this pricing signal relative to more immediate concerns about asset quality.

The headline finding across all three predictions is that the proportion of the failed bank's assets retained by the FDIC is the dominant quantitative lever. Because the retained asset share and the deposit premium are jointly determined at the auction, I instrument the realized toxic ratio with the pre-failure loan-loss provision (LLP) ratio, a balance-sheet measure determined before any bidding occurs. The instrumented toxic ratio raises bid premiums in the pricing channel, raises the probability of a P&A sale in the selection channel, and the raw provision ratio is the dominant driver of FDIC costs. A one percentage point increase in the loan-loss provision ratio raises the

cost to the Deposit Insurance Fund by approximately 86 basis points of total assets, making asset quality the single most economically significant determinant of fiscal exposure. Consistent with the theoretical framework, accelerated deposit runoff and high interest-rate sensitivity reduce bid premiums and increase resolution costs in the directions the model predicts, though with more attenuated statistical significance relative to the asset-quality channel. The most robust selection predictor is log asset size: larger institutions attract P&A acquirers at rates disproportionate to their deposit-franchise quality, a pattern consistent with the Too-Big-to-Fail dynamic embedded in the resolution process.

This paper contributes to several strands of the banking literature. First, it extends the early work on resolution costs by Bovenzi and Murton (1988) and James (1991) by providing a more granular decomposition of the cost structure using data that covers the post-GFC and 2023 crisis periods. While those studies established the determinants of aggregate losses, this paper isolates the specific pricing of the deposit franchise and the cost of asset retention. Second, it complements the recent literature on fire sales and bidder behavior by Granja, Matvos, and Seru (2017). While their work focuses on the geographic and portfolio fit of acquirers, this paper focuses on the explicit contract terms and the allocation of risk between the state and the private sector. Finally, it contributes to the literature on deposit franchise value and monetary policy transmission. By linking deposit betas to realized auction premiums, I provide a market-based validation of the theory that core deposits represent a valuable intangible asset even in the context of distress.

The construction of this dataset serves a broader purpose beyond the specific empirical tests presented here. By making these transaction-level details public, this paper opens the door for future research into the political economy of bank resolutions. The variation in loss-sharing agreements, the timing of closures, and the treatment of uninsured depositors across different regulatory regimes can now be scrutinized with greater precision. The data reveals a stark evolution in the protection of uninsured depositors where nearly all uninsured depositors were protected during the systemic events of 2008 and 2023 while non-crisis periods show a much higher incidence of losses. This cyclicality in depositor protection has profound implications for market discipline and moral hazard which are themes that are central to the ongoing debate over deposit insurance reform and the "Too-Big-To-Fail" doctrine.

The remainder of the paper is organized as follows. Section II provides the institutional background on the FDIC resolution process and reviews the relevant literature. Section III details the data collection methodology and describes the construction of the novel variables. Section IV presents the theoretical model of the failed bank auction. Section V reports the empirical results regarding the pricing, cost, and selection channels. Section VI concludes.

II. Institutional Details

This section provides an overview of the Federal Deposit Insurance Corporation's (FDIC) bank resolution process, including the typical timeline and the methods used to resolve failed banks. It

begins with an introduction to how resolutions unfold under the Federal Deposit Insurance Corporation Improvement Act of 1991 (FDICIA), outlining the key steps from regulatory intervention through closure and bidding. It then describes in detail each major resolution method employed by the FDIC: Purchase and Assumption transactions (P&As), Insured Deposit Transfers (IDTs), Insured Deposit Payoffs (POs), Purchase of Insured deposits (PIDs), and the use of a Deposit Insurance National Bank (DINB); explaining what each method entails, how it works, the cost implications for the FDIC and acquiring institutions, and the treatment of insured versus uninsured depositors. The section concludes with a brief literature review of empirical studies on United States bank failures and resolution costs, highlighting findings from academic research and FDIC analyses that provide context on the outcomes and costs associated with different resolution strategies.

A. S&L Crisis and the FDICIA

The Federal Deposit Insurance Corporation Improvement Act of 1991 (FDICIA) 102nd Congress of the United States (1991) is best understood as a legislative response to the intertwined thrift and commercial bank crises of the 1980s, and in particular to the policy failures revealed by the savings and loan (S&L) debacle. The S&L crisis exposed the fragility of a regulatory regime that combined generous, underpriced deposit insurance with rapid deregulation, weak supervision, and extensive regulatory forbearance. FDICIA's core provisions: prompt corrective action, least-cost resolution, and risk-based insurance premiums were explicitly designed to correct those incentive problems and to limit the fiscal and political costs of future failures Benston and Kaufman (1997).

A brief review of the S&L crisis highlights why such reforms became politically unavoidable. In 1980, roughly 4,000 federally insured S&Ls held portfolios concentrated in long-term, fixed-rate mortgages that became unsustainable when short-term interest rates spiked[1], leaving hundreds of institutions economically insolvent and ultimately generating resolution costs in excess of $150 billion, much of it borne by taxpayers (Curry and Shibut (2000); Federal Deposit Insurance Corporation and Resolution Trust Corporation (1998)). Deregulatory measures, notably the Depository Institutions Deregulation and Monetary Control Act of 1980 and the Garn-St Germain Act of 1982, removed interest rate ceilings, raised deposit insurance limits, expanded asset powers, and formalized capital forbearance, allowing weak thrifts to attract insured funds and "gamble for resurrection" while the Federal Savings and Loan Insurance Corporation (FSLIC) lacked the resources and political backing to close them promptly McCoy (2007). Congress eventually responded with the Financial Institutions Reform, Recovery, and Enforcement Act of 1989 (FIRREA), which abolished FSLIC, created the Resolution Trust Corporation to resolve failed thrifts, and transferred

[1] A key macroeconomic backdrop to the S&L crisis was the sharp rise in short-term interest rates associated with the Federal Reserve's disinflation campaign under Chair Paul Volcker. In October 1979, the Fed shifted its operating procedure to target nonborrowed reserves to bring down double-digit inflation; the resulting monetary tightening pushed the effective federal funds rate from roughly 10-11% in mid-1979 to peaks above 17% in late 1979-early 1980, and again to nearly 19-20% in 1981. From 1980 to 1982, the funds rate averaged in the low to mid-teens, while yields on short-term Treasury bills also moved into double-digit territory. This abrupt increase in funding costs, driven by anti-inflation policy rather than thrift-specific factors, devastated the traditional S&L business model of funding long-term, fixed-rate mortgages with short-term deposits.

thrift insurance to a new fund administered by the FDIC; however, the contemporaneous wave of commercial bank failures made clear that the broader design of deposit insurance and resolution policy required more fundamental reform.

Before FDICIA, the FDIC enjoyed broad discretion over failure resolution, choosing among deposit payoffs, purchase-and-assumption transactions, and open-bank assistance without a statutory least-cost test and with considerable latitude to protect uninsured depositors. High-profile interventions, such as the rescue of Continental Illinois in 1984, signaled a willingness to shield all depositors and many creditors of large banks, reinforcing expectations of "too-big-to-fail" protection and encouraging risk-taking Davison (1997). Supervisors also sometimes delayed closure, engaging in regulatory forbearance that allowed losses to grow before resolution Benston and Kaufman (1997).

FDICIA sought to harden these soft constraints; Prompt Corrective Action linked supervisory interventions to objective capital categories, aiming to force earlier, rules-based responses to deterioration and to limit forbearance. The least-cost resolution requirement constrained the FDIC's use of open-bank assistance and broad creditor protection, subject to a narrow systemic-risk exception. At the same time, risk-based insurance premiums and target reserve ratios shifted a greater share of the cost of failures onto insured institutions. Complementary provisions, such as more frequent examinations for weaker banks and limits on brokered deposits, reinforced these incentives. In aggregate, FDICIA represents a codified lesson of the S&L era: generous insurance and deregulation must be paired with credible, timely intervention and loss-allocation rules if public exposure to banking crises is to be contained.

B. FDIC Resolution Process

B.1. Timeline

Under FDICIA's Prompt Corrective Action framework, regulators follow a structured timeline when dealing with a failing bank. In general, a failing institution is closed by its chartering authority once it becomes insolvent or "critically undercapitalized" (typically defined as tangible equity $\leq 2\%$ of assets). At that point, the FDIC is appointed as receiver and must resolve the bank in a manner that is "least costly" to the deposit insurance fund, barring a systemic risk exception. The resolution timeline usually unfolds over about 90 days for a capital insolvency scenario, during which the FDIC prepares for an orderly failure resolution. Key steps in this process include:

1. **Regulatory Intervention Period**: Once the chartering agency determines a bank is likely to fail, the bank is given up to 90 days to attempt corrective actions or find a private-sector solution (recapitalization or a buyer) under FDICIA's prompt corrective action provisions. During this period, the FDIC begins contingency planning for resolution in case the bank cannot recover.
2. **Failure Determination and Receivership**: If the institution cannot be saved by the end of the 90-day window, the chartering authority (state or federal) formally closes the bank

and appoints the FDIC as receiver. This typically occurs once it is clear that no viable recapitalization or merger is forthcoming and the bank's capital is critically deficient.

3. **Marketing to Potential Acquirers**: The FDIC quickly moves to market the failing bank to potential bidders. Eligible bidder institutions are identified, usually well-capitalized banks in good regulatory standing capable of absorbing the failed bank. After signing confidentiality agreements, these qualified bidders are given access to due diligence information about the failing bank's assets, liabilities, and operations. The FDIC simultaneously develops a marketing strategy and resolution structure for the sale.
4. **Auction and Bid Submission**: The FDIC conducts a sealed-bid auction for the failed bank. Bidders typically submit bids comprising three components: (a) an offer (if any) to assume the failed bank's deposits, often with a deposit premium they are willing to pay for the deposit franchise, (b) an offer for some or all of the failed bank's assets, often at a discount to book value, and (c) an indication of whether the bid is for a whole-bank transaction: all deposits, including uninsured, or only the insured deposits. Bidders may submit multiple bid options[2]. In some cases, the FDIC may offer optional loss-sharing agreements or other structures to encourage bids on bad assets.
5. **Least Cost Bid Selection**: Once bids are received, the FDIC evaluates them using a proprietary Least Cost Test to determine which option would minimize the cost to the Deposit Insurance Fund (DIF). FDICIA mandates that the FDIC must choose the resolution alternative that is expected to be the least costly to the DIF, unless a systemic risk exception is invoked. In practice, this means that bids resulting in lower estimated losses to the FDIC, after accounting for assets recovered and any liabilities assumed by the acquirer, will be favored. The least-cost requirement also generally implies that uninsured depositors and other general creditors should absorb losses rather than be protected, unless protecting them is necessary to secure a bid or prevent systemic harm.
6. **Final Resolution**: Once the winning bid is selected, the FDIC prepares the legal closing documents, and the bank is closed. By the next business day, the FDIC carries out the resolution: either transferring operations to the acquiring institution or paying out deposits, depending on the method used. The entire resolution process from initial regulatory warnings to final closing usually occurs within about 90 days for capital failures. However, in liquidity crises, this timeline may be compressed.

B.2. Resolution Types

The FDIC employs several resolution methods to address failed banks, each with different implications for the continuity of banking services, FDIC costs, and outcomes for depositors. The primary methods include Purchase and Assumption transactions (P&As), Insured Deposit Transfers (IDTs), Insured Deposit Payoffs (POs), Purchase of Insured deposits (PIs), and, in rare cases, the creation of a Deposit Insurance National Bank (DINB). In all cases, the goal is to protect depositors

[2]For example, a bid assuming all deposits versus a bid assuming only insured deposits.

up to the insurance limit and dispose of the failed bank's assets and liabilities in an orderly way, while minimizing the FDIC's losses. The choice of resolution method depends on factors such as the bank's size and condition, bidders' interests, and the results of the least-cost test. Below, I describe each method in detail:

1. **Purchase and Assumption**: A P&A transaction is the most common resolution method, used in roughly 77.6% of bank failures historically. In a P&A, the "assuming institution" purchases some or all of the failed bank's assets and assumes some or all of the liabilities, especially the deposits. The FDIC acts as receiver to facilitate this transfer. P&A agreements are flexible and can be structured in various ways. In a whole-bank P&A, the acquirer assumes all deposits, including uninsured deposits, and takes on essentially all assets, even troubled loans, leaving very little with the FDIC. In other P&A variants, the acquirer may assume only the insured deposits, leaving uninsured depositors as creditors of the receivership, and may also purchase a portion of the assets. In any P&A structure, the FDIC typically sweetens the deal to ensure it meets the least-cost test: it will reimburse the acquirer for the difference between liabilities assumed and assets purchased by injecting cash or retaining bad assets, so that the acquiring bank does not take on a net loss at closing. Essentially, the acquirer buys the failed bank's good assets and assumes the deposits. At the same time, the FDIC absorbs the loss on bad assets by paying out receivership funds to balance the transaction. Often, the acquirer pays the FDIC a deposit premium, a cash payment in exchange for acquiring the deposit franchise, if the franchise has value, which offsets the FDIC's cost. P&A deals allow for continuity of banking services: the failed bank's branches typically reopen as branches of the acquiring institution immediately, and all insured depositors automatically become customers of the acquirer with uninterrupted access to their accounts. If the P&A is a whole-bank transaction, even uninsured depositors are fully protected, and the acquirer assumes their accounts. However, the FDICIA tightened the rules governing the protection of uninsured depositors. Since 1994, the FDIC has been prohibited from covering uninsured deposits in a resolution except when a systemic risk exemption is invoked. Thus, in modern P&A transactions under the least-cost mandate, uninsured deposits are assumed only by the acquirer if doing so does not increase the FDIC's cost or if regulators determine that not protecting them would threaten financial stability. Overall, P&A transactions tend to be less costly to the FDIC than straight liquidations because they preserve franchise value, the premium paid by an acquirer, and they help avoid asset fire-sales. The acquiring bank benefits by gaining new customers and assets, often at a discount or with loss protection. Indeed, during the Great Financial Crisis of 2008, many P&A deals included FDIC loss-sharing agreements (where the FDIC agreed to absorb a large percentage of future losses on certain loan portfolios), which further encouraged bidders to participate. These arrangements greatly limit the acquirer's risk, effectively providing a subsidy. Empirical evidence shows that acquirers in FDIC P&A auctions during the crisis realized significant gains. For example, one study found that winning bidders' stock prices jumped on the announcement of the

acquisition, reflecting the advantageous terms and FDIC support they receivedCowan and Salotti (2015).

2. **Insured Deposit Transfers**: An IDT is a resolution method typically used when no open bank is willing to acquire the entire failing institution, but the FDIC can at least arrange to transfer the insured deposit accounts to another bank. In an IDT, the FDIC transfers all insured deposits of the failed bank to a healthy bank, which agrees to act as the paying agent for those deposits. The receiving bank does not purchase the failed bank as a going concern; rather, it simply opens accounts on its books for the failed bank's insured depositors so that those customers can quickly access their insured funds, usually at full-service bank branches in their community Barth and Prabha (2013). The FDIC, in its receivership capacity, provides cash (or quickly sellable assets) equal to the total insured deposit amounts to the acquiring institution to cover the transferred deposits. Essentially, this is a facilitated payout: the insured depositors receive their money by seamlessly becoming depositors of the agent bank, rather than receiving a check from the FDIC. Uninsured depositors in this scenario are not transferred; they remain creditors of the receivership for the amount of their uninsured balances, to be paid out from asset liquidation recoveries (if any) over time. The agent bank typically receives a fee or other considerations for assisting, or sometimes pays a modest premium if the deposit franchise is attractive. However, this method is usually employed when franchise value is low.

   The IDT method was more common in earlier decades and has been used rarely since the mid-1990s; in fact, the FDIC has not employed an IDT since 1995. It is considered similar to a direct payoff in that the failed bank is liquidated, with no ongoing operations sold. However, it has the advantage of giving insured depositors uninterrupted access to their funds through a new bank, thereby avoiding service disruptions that a payoff by mailed checks would entail. For the FDIC, the cost of an IDT is comparable to an insured deposit payoff. The FDIC still bears the loss for the difference between the failed bank's assets and insured liabilities. However, it may reduce resolution costs if an acquiring institution pays a small premium for the deposits or if having a ready agent bank streamlines the claims process. From the depositor's perspective, insured customers are made whole and can quickly access their funds at the acquiring bank. In contrast, uninsured depositors incur losses, receiving only receivership certificates and, if asset recoveries permit, eventual dividends.

3. **Insured Deposit Payoffs**: A deposit payoff (PO), also known as a deposit payout or liquidation, is the resolution method of last resort when no assuming institution can be found. In a payoff, the FDIC, as receiver, pays insured depositors directly the amount of their insured funds and winds down the failed bank's operations. The bank is closed and liquidated: all assets of the failed bank are retained by the FDIC receivership for sale over time, and the proceeds are used to reimburse the FDIC for insured payouts and to pay any remaining claims by uninsured depositors or other creditors, in order of priority. Uninsured depositors in a payoff do not get an assuming bank; instead, they receive a receivership certificate and may

eventually recover a portion of their uninsured balances through liquidation dividends if the asset sales yield sufficient funds. In most cases, uninsured depositors suffer losses (a partial recovery) because failed bank assets often liquidate for less than the total liabilities. From the FDIC's perspective, a straight payoff is usually the most costly resolution method, because it forgoes any franchise value. The FDIC must immediately draw on the insurance fund to pay all insured deposits, and then bear the risk and expense of liquidating the assets over time. Payoffs tend to result in higher FDIC losses as a percentage of the failed bank's assets compared to P&A transactions. POs also disrupt banking services: the failed bank does not reopen, and customers must wait to receive their insured funds (often by check or electronic transfer) and then find a new bank on their own. Because of these drawbacks, deposit payoffs have been relatively infrequently used in roughly 7% of failures in past decades and generally only occur when no acceptable bid is offered for a P&A or transfer. Notably, even in a payoff, the FDIC may use an IDT or DINB as a vehicle to facilitate the payout, as described next, to reduce disruption.

4. **Purchase of Insured Deposits**: A PI resolution is essentially a hybrid of a P&A and an IDT/PO. In a PI transaction, an acquiring bank agrees to purchase only the insured deposit liabilities of the failed bank, without taking on the uninsured deposits. The acquirer usually also takes a portion of the assets equal to those deposits, often mostly cash and investment securities, or the FDIC may provide cash to fund the assumed deposits. In effect, the acquirer bank is stepping into the FDIC's shoes to pay insured depositors, but through a P&A agreement for the insured deposit portion. The acquirer may pay the FDIC a small premium for the insured deposit franchise since it gains a base of new customer accounts, which can reduce the FDIC's net cost compared to a payoff. However, the acquirer leaves behind all uninsured deposits and any unwanted assets to be handled by the FDIC's receivership. From the depositor's standpoint, insured depositors are protected and become customers of the acquiring bank, just as in a P&A or IDT, they retain access to their accounts, now at the new bank. Uninsured depositors, by contrast, are not assumed by the acquirer; they suffer an interruption and must seek recovery from the receivership for the uninsured portion of their funds, likely recovering only a fraction after asset liquidation. PI transactions can be viewed as a partial P&A: they preserve some insured deposit relationships but not the entire franchise. They have been used in cases where an acquirer is willing to take on the deposit base but perhaps not the riskier assets or the uninsured liabilities. This method was employed in a minority of failures, around 6% historically. The cost to the FDIC in a PI is typically lower than a full payoff because the FDIC might receive a premium and avoid the logistical costs of directly paying insured deposits. Still, because the FDIC must cover losses on assets not taken by the acquirer and handle the uninsured claims, the DIF will incur losses on a PI, albeit hopefully less than in a straight liquidation. Acquiring banks in PI deals gain insured deposit accounts without exposing themselves to the failed bank's credit risks since they can cherry-pick assets or take none at all beyond cash, making it an attractive

proposition for the acquirer with minimal risk.

5. **Deposit Insurance National Bank**: A DINB is a temporary national bank chartered by the FDIC to maintain banking services for depositors of a failed bank when no immediate acquirer is found and a straight payoff would be disruptive. In this resolution, the FDIC itself obtains a new national bank charter from the Office of the Comptroller of the Currency and operates the DINB solely to provide depositors with access to their insured funds while they transition to other banks. The DINB takes over the insured deposits of the failed bank; depositors can access their insured balances by doing business with the DINB, which might operate out of the failed bank's existing offices for a limited period. By law, a DINB can operate for up to 2 years (with possible extensions), but in practice, the FDIC usually keeps a DINB open for only a short time (a few weeks to a few months). During that time, depositors are expected to move their funds to other institutions. The DINB does not make loans or conduct a full banking business; it exists solely to facilitate the smooth payout of insured deposits. It has no capitalization requirements because the FDIC backs it. The DINB does not assume all uninsured deposits of the failed bank; initially, those uninsured depositors would receive receivership certificates and stand to recover funds later as assets are sold. In some cases, if a buyer for the whole bank emerges shortly after failure, the DINB can be bridged into a purchase transaction. For example, in the March 2023 failure of Silicon Valley Bank, the FDIC first created a DINB to pay insured depositors and was preparing to handle uninsured claims through receivership. However, within days, it transferred all deposits, including uninsured ones, into a newly created bridge bank after invoking a systemic risk exception. DINBs are exceedingly rare; they have accounted for less than 1% of failures because the FDIC usually prefers to either find an acquirer or, if possible, pay off deposits directly. However, the DINB tool is available for situations where the sudden loss of banking services would significantly harm the community or where an orderly payoff requires maintaining an interim bank. The FDIC's cost for a DINB is similar to a payoff; it still ultimately pays out insured deposits and bears the costs of asset liquidation and the administrative expenses of operating the temporary bank. Nevertheless, the benefit is that it gives depositors time to make alternative arrangements and can preserve a semblance of normalcy in the immediate aftermath of a bank failure, potentially preventing local panic.

In summary, the FDIC's resolution toolkit under FDICIA is designed to protect insured depositors while minimizing the insurance fund's losses. Purchase and Assumption transactions are the preferred method in most cases, as they maintain continuity and often yield the lowest cost by leveraging an acquiring institution. Insured deposit transfers and DINBs are fallbacks to ensure insured depositors are made whole when no outright buyer is available. A deposit payoff is the least favored due to higher costs and disruption, and is used only when necessary. The treatment of uninsured depositors depends on the method and the least-cost calculus: post-FDICIA, uninsured depositors are generally at risk of loss unless the FDIC can arrange a whole-bank acquisition or invoke a systemic risk exception to protect them.

C. *Related Literature*

The empirical literature on bank failures and FDIC resolutions is large and spans several distinct strands: early work that measures resolution losses and their determinants at the bank level; studies that embed these experiences in broader evaluations of deposit insurance and insolvency regimes; and more recent research that exploits detailed transaction-level data from FDIC auctions, often focusing on bidder behavior, acquirer outcomes, and the incidence of losses across claimants. This subsection briefly reviews the main contributions in each strand, with an emphasis on U.S. evidence and on studies that explicitly analyze the FDIC's choice of resolution method and its implications for the deposit insurance fund and depositors.

Early empirical work by FDIC staff and affiliated researchers documents basic facts about resolution costs across failure episodes and resolution methods. Bovenzi and Murton (1988) and Bovenzi and Muldoon (1990) use FDIC Banking Review data to estimate the average costs of alternative resolution methods, discussed in Section II.B.2, and show that P&A transactions generally reduce losses relative to straight liquidations when suitable acquirers exist. James (1991) measures realized losses as the difference between book liabilities and discounted recovery values of failed bank assets, documenting that losses are strongly related to pre-failure asset quality, capital ratios, and local economic conditions. Brown and Epstein (1992) updates the FDIC's historical loss model, while subsequent work by McDill (2004) links resolution costs to the business cycle and macroeconomic stress. Parallel work on the thrift sector, such as the analysis of the savings and loan crisis in Curry and Shibut (2000), quantifies the aggregate fiscal burden of large failure waves and emphasizes that taxpayers ultimately bore a substantial share of losses via the deposit insurance funds. These studies establish the now-standard definition of "resolution cost" as the shortfall between the failed bank's liabilities and the discounted market value of its assets net of resolution expenses, and they highlight systematic differences in this cost across resolution technologies and institutional environments.

A second strand embeds these experiences within a broader evaluation of deposit insurance design and the legal framework for bank resolution. FDIC historical volumes, notably Managing the Crisis: The FDIC and RTC Experience, 1980-1994 Federal Deposit Insurance Corporation and Resolution Trust Corporation (1998), provide a narrative and quantitative account of the 1980s banking and thrift crises, including the extensive use of P&A transactions, bridge banks, and open-bank assistance, as well as the evolution of FDIC practice in response to large-scale failures. Benston and Kaufman (1997) interpret the FDICIA as a response to moral hazard problems created by mispriced deposit insurance and forbearance during the 1980s, and they emphasize two innovations with direct implications for resolution policy: prompt corrective action and the statutory least-cost resolution requirement. Bliss and Kaufman (2006) compare U.S. corporate bankruptcy and bank receivership regimes, stressing that bank resolution is an administrative process run by the FDIC under the Federal Deposit Insurance Act rather than a judicial Chapter 11 proceeding, with different objectives (maintaining confidence and limiting systemic risk), different priority rules (including depositor preference), and much faster timelines for closing, valuing, and reprivatizing

failed institutions. Related work on depositor preference, secondary deposit insurance, and changes in creditor priority (e.g., Hirschhorn and Zervos (1990); Marino and Bennett (1999); Osterberg (1996)) clarifies how statutory and contractual priority rules govern the distribution of resolution losses across insured and uninsured depositors, general creditors, and the FDIC.

More recent research exploits detailed FDIC data to dissect the components of resolution cost and to study the auction process used to sell failed banks. Bennett and Unal (2015) construct a reconciled time series of cash flows for more than 1,200 FDIC receiverships between 1986 and 2007, decomposing total resolution costs into losses on assets, direct receivership expenses, and indirect FDIC overhead, and documenting how these components vary with resolution method, bank size, legislative period (pre-FIRREA, FIRREA[3], FDICIA), and fraud. They show, for example, that resolution costs as a share of assets decline with bank size and are lower in the post-FDICIA period, consistent with earlier closures and the operation of the least-cost test. Complementary FDIC publications, such as Failed Bank Cost Analysis: 1986–2002 Corporation (2003) and subsequent staff studies, extend this line of work by formalizing FDIC cost estimation methodologies and relating realized costs to ex-ante projections used in the least-cost decision. A parallel literature, starting with Granja et al. (2017), uses FDIC auction data to study which potential acquirers win failed bank auctions and how bidder characteristics—geographic proximity, portfolio similarity, and capitalization—shape both allocation and FDIC losses, interpreting these findings through a fire-sale framework.

Finally, there is a growing body of event-study and contract-level work on FDIC-assisted transactions during the 2008-2013 crisis and its aftermath. Cowan and Salotti (2015) analyzes stock returns for winning bidders in FDIC-assisted acquisitions and shows that acquirers earn positive abnormal returns that are increasing in the generosity of FDIC guarantees, indicating that resolution design redistributes surplus between the FDIC and acquiring banks. Heitz (2022) focuses on loss-sharing agreements in failed-bank sales and finds that acquiring banks entering into loss-sharing arrangements experience significantly lower abnormal returns than otherwise similar auction participants, implying that loss-sharing shifts more downside risk back to acquirers and away from the FDIC. Recent work also examines externalities on rival and losing bidders (e.g., Molyneux and Zhou (2022)) and compares FDIC-assisted acquisitions to open-market mergers in terms of long-run profitability and risk (Cowan, Salotti, and Schenck (2022); Spokeviciute, Jahanshahloo, Keasey, and Vallascas (2025)). FDIC staff studies contrast the 1980-1994 SNL and 2008 GFC crises, emphasizing that, although far fewer banks failed in the latter episode, the average size of failures and the complexity of resolution strategies, including greater use of loss sharing and structured transactions, were much larger.

[3]The Financial Institutions Reform, Recovery, and Enforcement Act of 1989

III. Press Release Data

The empirical analysis relies on a novel dataset constructed by merging administrative failure records with hand-collected details from the FDIC's failure-specific press releases. The foundational population is drawn from the FDIC's BankFind Suite: Bank Failures and Assistance Data, which provides the universe of federally insured depository institution failures. This base file is augmented with granular transaction-level data extracted directly from the FDIC's press release archive and matched to quarterly Call Report data from Drechsler, Savov, and Schnabl (2017, 2021). The final sample covers the period from January 21, 1983 to October 18, 2024 and contains 1,860 observations. The panel consists primarily of commercial banks, national and state-chartered, and 68 savings banks.

This section has two purposes. First, it documents the construction and coverage of the hand-collected dataset. Second, it presents the descriptive cross-sectional and time-series variation in the six variables that the structural model in Section IV treats as primitives or outcomes: deposit runoff, the deposit beta of Drechsler et al. (2021), the loan-loss-provision ratio, the uninsured share, the realized retained-asset share, and the auction outcomes. Taken together, Tables I and II and the accompanying figures show that there is substantial cross-sectional dispersion in each primitive, and that simple bivariate patterns in the data already motivate the structural links that the model later formalizes.

A. Construction and Coverage

For each failed bank, the FDIC issues a standardized public announcement that describes the closing authority, the assuming institution, the timing of branch reopening, and the key balance-sheet and resolution terms. An example of the common template appears in Appendix B. From each release I code the volume of assets and deposits at the moment of resolution, the total number of deposit accounts, the subset of accounts and deposit balances above the \$100,000 (later \$250,000) insurance limit, the premium paid by the acquirer on deposits and on assets, the volume of assets retained by the FDIC in receivership, any cash advance from the FDIC to the acquirer, the notional amount of assets covered by a loss-sharing agreement, the FDIC's share of those losses, and the FDIC's projected resolution cost. A full list of the variables contained in the data set can be found in Table A. These variables are not available in standard FDIC statistical publications or call report data but are central to allocating resolution outcomes across acquirers, the DIF, and uninsured depositors. The Call Report match, following Bennett and Unal (2015), provides the pre-failure (QBF) balance sheet, from which I construct the loan-loss-provision ratio and the bank's size.

Table I documents the coverage of all variables used in the paper. Four remarks are in order. First, QBF and final balance-sheet items are available for essentially the full sample, as are the core auction outcomes: premiums on deposits and assets, retained assets, cost, and uninsured-depositor recovery. Second, the loan-loss provision from the Call Report is available for 1,785 of 1,860 failures.

Data Field	Count	% Missing
Panel A: Balance Sheet Inputs		
QBF Assets	1,860	0.00
QBF Deposits	1,860	0.00
Final Assets	1,860	0.00
Final Deposits	1,860	0.00
Deposit Runoff	1,858	0.11
Log Asset Size	1,860	0.00
Panel B: Theoretical Primitives		
LLP Ratio	1,785	4.03
Deposit Beta	373	79.95
Expense Beta	373	79.95
Uninsured Share	414	77.74
Core Deposit Share	414	77.74
Panel C: Auction Outcomes		
Recovery Category	1,860	0.00
Recovery %	1,860	0.00
Premium on Deposits	1,860	0.00
Premium on Assets	1,856	0.22
Retained-Asset Share	1,860	0.00
FDIC Advance	1,860	0.00
Loan-Loss Provision (press release)	1,860	0.00
Loan-Loss FDIC Share	7	99.62
Cost	1,860	0.00

Table I: Coverage and missingness of variables used in the paper ($N = 1{,}860$ failures). Panel A lists balance-sheet inputs from call reports and press releases. Panel B lists the theoretical primitives used as regressors. Panel C lists auction outcomes from the press-release dataset.

Third, the matched Drechsler et al. (2021) deposit betas cover 373 failures, concentrated in post-2000 closures. Fourth, the number and balance of accounts above the insurance limit are available for only 370-414 failures, primarily during the 1980s, when the FDIC reported this granularity in its releases. I therefore treat uninsured exposure as a restricted-sample covariate and flag the sample restriction explicitly in the regression tables in Section V.

B. Failure Landscape

Figure 1 in Section I traces the annual count and resolution-method mix of failures over the sample. Three features of that figure are worth carrying into the rest of this section. First, failures cluster in three waves: the S&L crisis of 1986-1995, the Global Financial Crisis (GFC) of 2008-2013, and the small but large-asset 2023 regional banking episode. Second, the share of resolutions structured as P&As rises from roughly three-quarters in the 1980s to nearly one in the Dodd-Frank period, with pure payoffs and insured deposit transfers becoming rare. Third, as noted above, average failed-bank asset size rises by nearly two orders of magnitude across the sample. However, medians remain in the community-bank range, a pattern already documented in Curry and Shibut

(2000) and Bennett and Unal (2015).

Figure 2 adds a spatial dimension by mapping the location of each failed bank by regulatory regime. Failures cluster persistently in Texas and the broader Southwest, the agricultural Midwest, California, and parts of the Northeast, with the concentration shifting across regions as local real estate and commodity cycles evolve. Many S&L-era failures involve small community banks in rural or semi-rural counties with limited branch competition. Morgan, Pinkovskiy, and Yang (2016) shows that post-GFC branch closures are associated with the emergence of "banking deserts" in which households and small firms face reduced access to basic deposit and credit services. The spatial pattern in Figure 2 is consistent with that concern. It underscores that resolution policy is not only about minimizing DIF losses but also about managing localized disruptions to relationship banking.

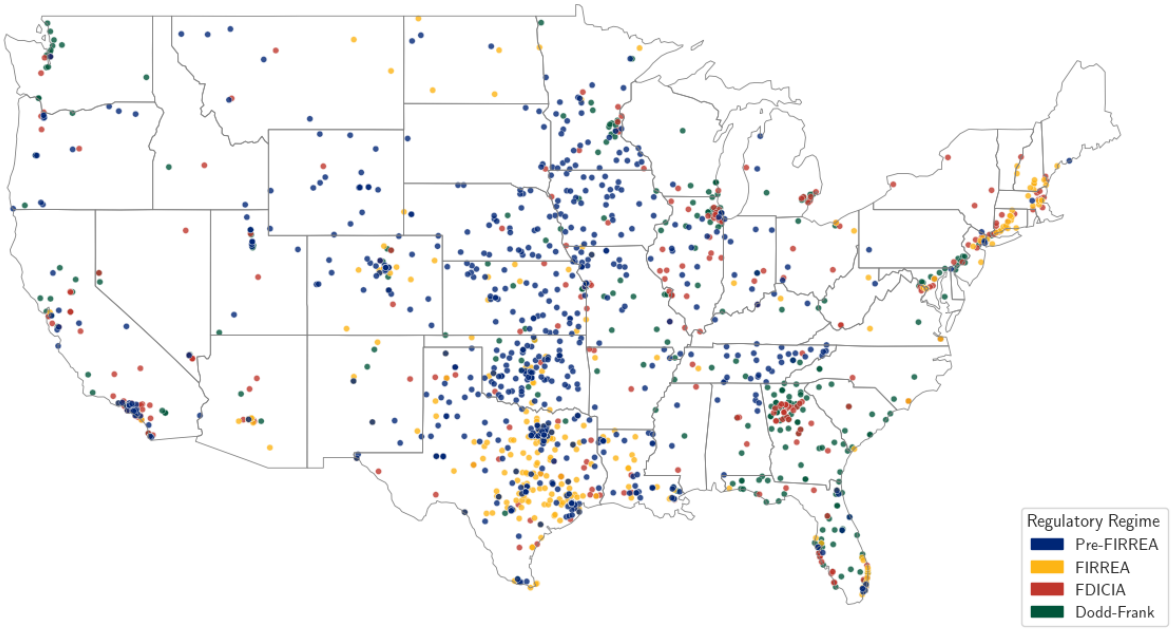


Figure 2. Locations of failed banks by regulatory regime. Each point is a failed institution colored by regime: Pre-FIRREA, FIRREA, FDICIA/PCA, or Dodd-Frank.

C. Balance Sheet Primatives

The model in Section IV treats four bank-level objects as exogenous primitives that determine franchise value and the equilibrium bid: the deposit attrition hazard λ, the deposit rate sensitivity β, the pre-resolution loan-loss provision ratio that proxies for asset toxicity γ, and the uninsured share of deposits. This subsection presents the cross-sectional variation in each.

The model's attrition hazard λ is not directly observable, but failing banks approach resolution over a short window of days to weeks. I use the log change in reported deposits between the last Call Report and the resolution date as a one-shot realization of cumulative attrition up to $t = 0$,

$$\text{Deposit Runoff}_i := \ln\left(\frac{\text{Final Deposits}_i}{\text{QBF Deposits}_i}\right), \tag{1}$$

where $t = -1$ is the last QBF Call Report and $t = 0$ is the resolution date. The timing convention matches the model.

Figure 3 plots the distribution of Deposit Runoff$_i$ for the full sample alongside the subsamples of failures that occurred during the three crisis windows. The full-sample distribution is centered near zero, with a pronounced left tail: a nontrivial share of banks lose more than 20 log points of deposits in the quarters leading up to closure. The 2023 subsample shifts the entire distribution to the left, consistent with the rapid deposit outflows documented in the March 2023 regional episode, while the S&L and GFC subsamples display heavier left tails than the full sample but retain substantial mass near zero. This variation is what identifies the coefficient on Deposit Runoff in the deposit-premium regressions of Section V.

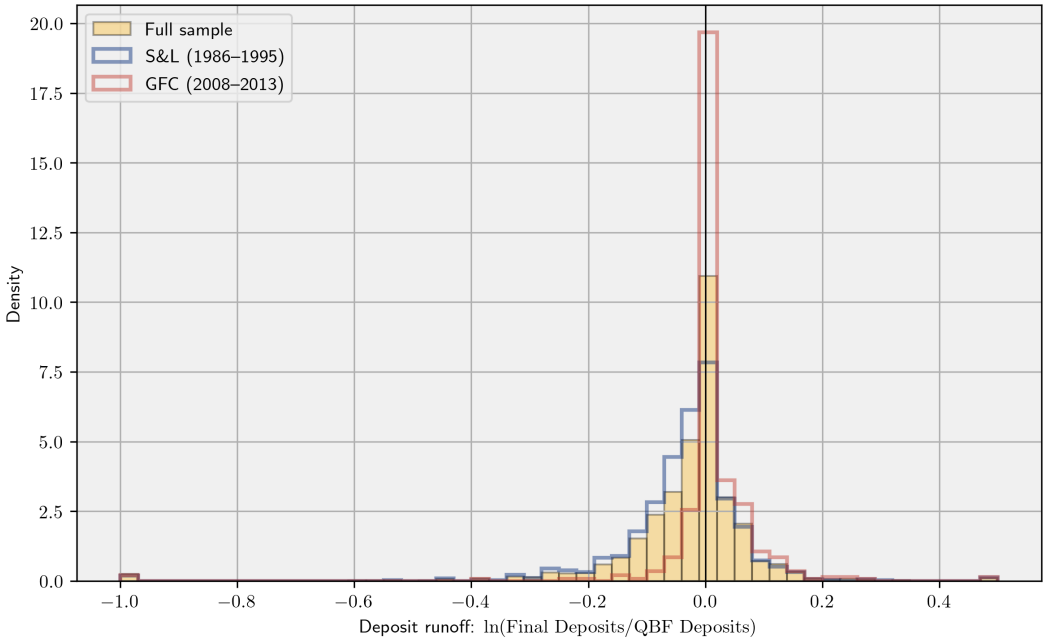


Figure 3. Deposit runoff distribution. Histogram of ln(Final Deposits/QBF Deposits) for the full sample, overlaid with density step plots for the S&L (1986-1995), GFC (2008-2013), and 2023 subsamples. Runoff values are clipped at $[-1.0, 0.5]$ for display.

Drechsler et al. (2021) estimate bank-level sensitivities of deposit rates to short-term policy rates and argue that low-beta deposits are the core source of franchise value because they insulate net interest income from rate shocks. I match each failed bank to the deposit beta and interest-expense beta posted on Schnabl's website; the match rate is 373 banks, concentrated in the post-2000 period. Figure 4 plots the log deposit premium from the FDIC auction against the matched deposit beta. The raw scatter is noisy, but the decile means on the right panel are monotonically decreasing: acquirers pay systematically higher premiums for banks whose deposits reprice slowly. This bivariate pattern is the empirical counterpart of the mapping from β to franchise value developed

in Section IV and is what identifies the coefficient on deposit beta in the premium regressions.

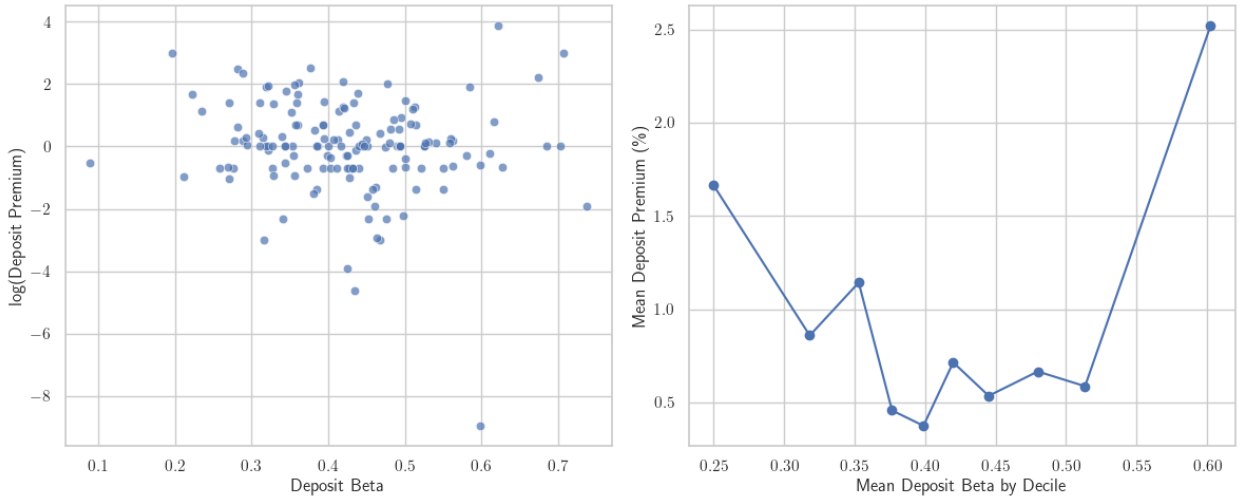


Figure 4. Deposit premium versus deposit beta. Left: scatter of log deposit premiums against bank-level deposit betas matched from Drechsler et al. (2021). Right: mean deposit premium by beta decile.

The model's asset-valuation wedge γ is proxied by the ratio of loan-loss provisions to QBF assets,

$$\text{LLP}_i := \frac{\text{Loan-Loss Provision}_i}{\text{QBF Assets}_i}, \tag{2}$$

computed from the last Call Report before closure and winsorized at the 1% and 99% tails. The identification logic is straightforward. Loan-loss provisions are an accounting recognition of expected credit losses that is booked before the FDIC opens the auction. They are therefore predetermined with respect to the auction's equilibrium outcome and can be treated as exogenous in regressions of cost and the retained-asset share, in the spirit of James (1991), who shows that realized resolution losses are strongly related to pre-failure asset quality.

Figure 5 plots the cross-sectional distribution of the LLP ratio. The distribution has substantial right skewness: the sample mean is 1.8%, but the 90th percentile exceeds 5%.

Figure 6 shows the first-stage relationship that motivates the use of LLP_i as an instrument for the realized retained-asset share τ_i in the cost regressions. Restricting to P&A transactions, I bin the sample into 20 quantile bins of the LLP ratio and plot the mean retained share in each bin, together with the OLS fit on the underlying observations. The slope is 1.18, so a one-percentage-point increase in the pre-failure LLP ratio is associated with a 1.2-percentage-point increase in the realized retained share. This pattern is consistent with the least-cost logic: banks with more troubled assets enter the auction with a larger share of assets that acquirers are unwilling to bid for at book value, and the FDIC retains the residual in receivership. It is the identifying variation behind the asset-side instrumentation in the analysis of Section V.

For the subsample of 370-414 failures in which the FDIC release reports over-the-limit balances, I construct the uninsured share as

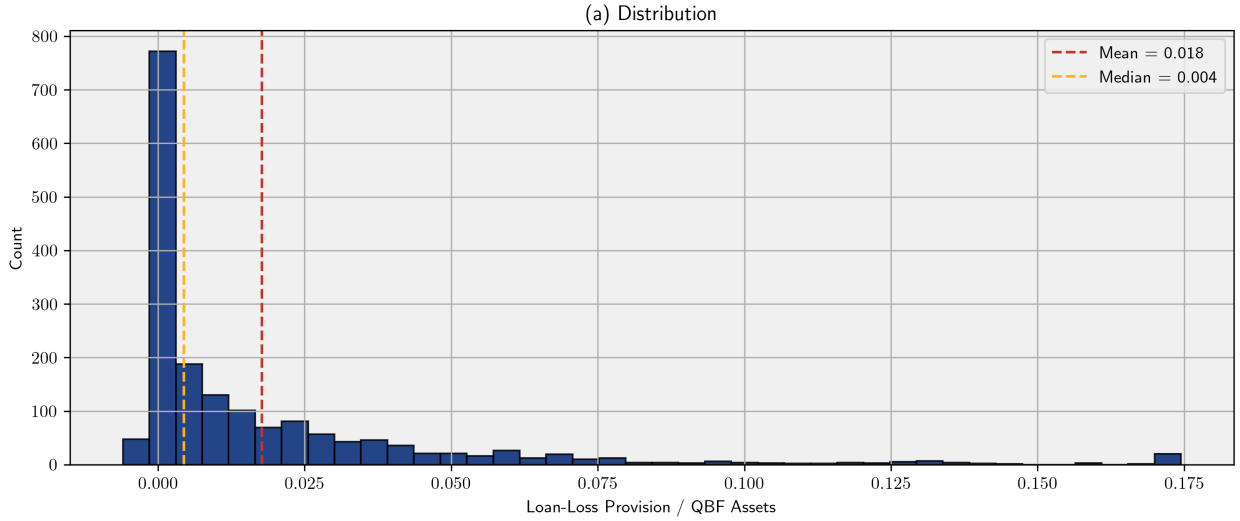


Figure 5. Loan-loss-provision ratio. Panel (a): histogram of LLP scaled by QBF assets, winsorized at 1% and 99%, with the sample mean and median. Panel (b): annual mean and interquartile band; shaded regions mark the S&L, GFC, and 2023 crisis windows.

$$\text{Unins}_i = \frac{\text{Amount Over Limit}_i}{\text{Final Deposits}_i} \tag{3}$$

and the complementary core-deposit share as $1 - \text{Unins}_i$. Coverage is heavily tilted toward the 1980s and early 1990s, and I flag this restriction whenever these variables enter the regression analysis. In line with the classical bank-run logic of Diamond and Dybvig (1983), uninsured balances are the claim most exposed to loss in a payoff, and variation in the uninsured share across failures is a natural shifter of the option value of running in resolution.

	N	Mean	Std. Dev.	p10	p25	Median	p75	p90
Deposit Runoff	1,858	-0.040	0.288	-0.123	-0.057	-0.002	0.014	0.058
Deposit Beta	373	0.416	0.100	0.297	0.354	0.408	0.481	0.540
Expense Beta	373	0.373	0.093	0.261	0.316	0.362	0.427	0.486
LLP Ratio (winsorized)	1,785	0.018	0.031	0.000	0.000	0.004	0.023	0.052
Retained-Asset Share	1,860	0.307	0.321	0.000	0.000	0.228	0.544	0.794
Uninsured Share	414	0.713	13.843	0.001	0.004	0.015	0.043	0.075

Table II: Summary statistics for the theoretical primitives used in the empirical analysis. Deposit Runoff is the log change in deposits between the last call report and the resolution date. Deposit and Expense Betas are matched from bank-level estimates following Drechsler et al. (2021). LLP Ratio is the pre-failure loan-loss provision scaled by QBF assets, winsorized at the 1% and 99% tails. Retained-Asset Share is the ratio of assets retained by the FDIC to QBF assets. Uninsured Share is deposits above the insurance limit as a fraction of final deposits.

Table II collects the cross-sectional moments of the six primitives. The interquartile range of deposit runoff indicates that half of failed banks experience at most modest attrition in the quarter before closure, but a meaningful minority experience 5-15% deposit attrition. Deposit and expense betas range from roughly 0.2 to 0.7 across the matched subsample, so a wide cross-section identifies

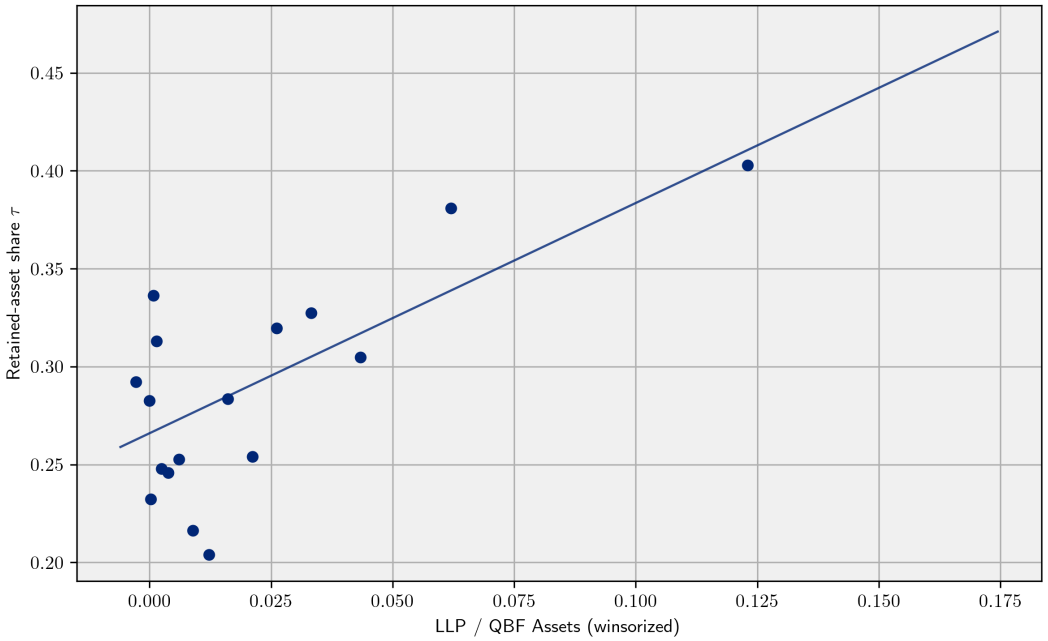


Figure 6. Realized retained-asset share vs. pre-failure LLP ratio, P&A sample. Points are means within 20 quantile bins of LLP/QBF Assets; the line is the OLS fit on the underlying observations. The slope and HC1 standard error are reported in the title.

the model's rate sensitivity parameter. The retained-asset share has a mean of 0.31 but a median of 0.23, with 30% of failures recording $\tau_i = 0$, a whole-bank P&A or straight payoff.

D. Resolution Outcomes

I now turn to the outcome variables that the model's equilibrium bid determines: the premiums paid on deposits and assets, the realized retained-asset share τ, the projected cost to the DIF, and the recovery paid to uninsured depositors. Before documenting the unconditional shape of these outcomes, I first describe how failures and resolution methods vary across the legislative and macro-financial environment in which they occur. The institutional regime fixes the menu of resolution technologies available to the FDIC and the statutory tests the agency must satisfy, and the macro-financial regime shapes both the size and the asset composition of the banks that resolve. Both margins matter for interpreting the pricing moments reported below.

Table III summarizes the distribution of failures across legislative regime (Panel A) and macro-financial environment (Panel B). Two shifts are first-order for interpreting the auction outcomes. First, mean QBF assets rise by nearly two orders of magnitude across regimes, from $43 million under pre-FIRREA to $1.9 billion under Dodd-Frank, with a further order-of-magnitude jump to $106 billion for the five banks in the 2023 crisis subsample. The contrast between the exploding mean and the community-bank-sized median documented in Bennett and Unal (2015) is therefore entirely driven by a small number of large failures concentrated in the FDICIA and Dodd-Frank windows. Second, the composition of resolution methods changes discontinuously at each legislative

	# Failures	Assets	Deposits	P&A	IDT	PO	PI	DINB
Panel A: Regulatory Periods								
Pre-FIRREA (1983–1988)	714	42.53	38.62	0.749	0.153	0.098	0.000	0.000
FIRREA (1989–1991)	420	151.81	132.47	0.840	0.110	0.050	0.000	0.000
FDICIA (1992–2009)	363	1,443.62	1,015.84	0.573	0.047	0.080	0.292	0.008
Dodd-Frank (2010–2024)	363	1,906.31	1,593.72	0.956	0.000	0.022	0.008	0.014
Panel B: Crisis Periods								
S&L Crisis (1986–1995)	1,052	99.439	89.346	0.718	0.144	0.072	0.067	0.000
Global Financial Crisis (2008–2013)	471	1,340.25	977.03	0.930	0.000	0.034	0.019	0.017
2023 Regional Crisis (2023)	5	106,445.68	88,122.01	1.000	0.000	0.000	0.000	0.000
No Crisis (1983–1985, 1996–2007, 2013–2022, 2024)	332	126.68	107.53	0.738	0.063	0.108	0.090	0.000

Table III: Summary of failed banks by regulatory regime (Panel A) and crisis environment (Panel B). Columns report counts of failures, mean quarter-before assets and deposits (in millions of dollars), and the share of resolutions by type: purchase and assumption (P&A), insured deposit transfer (IDT), payout (PO), purchase of insured deposits (PI), and deposit payoff with DINB bridge bank (DINB).

boundary. In the pre-FIRREA period, 15% of resolutions are structured as insured deposit transfers and roughly 10% as straight payoffs, so that, by construction, more than one in four failures impose losses on uninsured depositors. Under FIRREA, the IDT and PO shares fall. Under FDICIA, the purchase of insured deposits (PI) appears and accounts for 29% of resolutions, reflecting the least-cost logic of Benston and Kaufman (1997): the PI preserves the insured-deposit franchise for an acquirer while leaving uninsured depositors and distressed assets in the receivership. The Dodd-Frank regime then returns to a near-exclusive reliance on whole-bank P&A transactions, with 96% of failures resolved by a single acquiring institution, consistent with the evidence in Spokeviciute et al. (2025) that the FDIC has increasingly favored franchise-preserving resolutions over the last three decades.

Panel B of Table III cuts the sample by macro-financial environment. Three features are worth emphasizing. First, the S&L crisis accounts for 1,052 of the 1,860 failures in the sample but for only a small share of aggregate failed-bank assets, confirming the "many small failures" character of the episode narrated in Federal Deposit Insurance Corporation and Resolution Trust Corporation (1998) and quantified in Curry and Shibut (2000). Second, the GFC and the 2023 crisis invert this pattern: 471 and 5 failures, respectively, but mean assets an order of magnitude (GFC) and three

orders of magnitude (2023) larger than in the S&L episode. The 2023 subsample is essentially the mass of the large regional banks resolved under systemic-risk exceptions. Third, the P&A share is monotone in crisis severity: 72% in the S&L crisis, 93% in the GFC, and 100% in 2023. This pattern is consistent with the view in Granja et al. (2017) that, conditional on a sufficiently deep pool of well-capitalized bidders, whole-bank P&As dominate the least-cost calculus even under significant asset-quality stress, and with the Cowan and Salotti (2015) evidence that FDIC loss-sharing and cash advances were used aggressively during the GFC to maintain this preference for franchise-preserving resolutions. The "no crisis" subsample, by contrast, looks more like a low-intensity extension of the S&L pattern: a modestly larger average bank, a slightly lower P&A share, and a small but persistent use of IDTs and POs for the smallest and least attractive institutions.

The heterogeneity across the rows of Table III is the institutional backdrop against which the pricing moments below are formed. In particular, (i) the asset-premium, deposit-premium and cost distributions pool banks resolved under four different statutory regimes, and (ii) the unconditional means are heavily influenced by the post-FDICIA era, in which the P&A has become the dominant technology and in which resolution losses are allocated between the FDIC, acquirers, and uninsured depositors through a least-cost-scored auction rather than a discretionary negotiation.

	N	Mean	Std. Dev.	Min	q25%	Median	q75%	Max
Cost / QBF Assets	1,860	0.230	0.138	0.000	0.132	0.218	0.312	0.937
Premium on Deposits (%)	1,860	1.053	2.358	-1.500	0.000	0.075	1.160	48.146
Premium on Assets (%)	1,856	-2.010	6.406	-78.226	0.000	0.000	0.000	5.791
Retained-Asset Share τ	1,860	0.307	0.321	0.000	0.000	0.228	0.544	3.234
Uninsured Deposit Recovery (%)	1,860	90.935	37.843	0.000	100	100	100	100
FDIC Advance ($m)	1,860	9.459	33.623	0.000	0.000	0.000	6.425	410.500

Table IV: Summary statistics for FDIC resolution outcomes. Cost is the projected DIF loss scaled by QBF assets. Premiums are reported as a percentage of deposits (assets). The retained-asset share is the ratio of retained assets to QBF assets. Recovery is the realized payoff to uninsured depositors.

Three moments in Table IV summarize the pricing of failed banks. First, the Cost/Assets ratio has a right tail but is tightly bunched near 22%, close to the figure cited in Bennett and Unal (2015) for the 1986-2007 period. Second, the average premium paid on assets is near zero and strictly negative when concentrated discounts are included. In contrast, the premium on deposits is strictly positive on average, reinforcing the observation of James (1991) that value in a P&A is a deposit-side phenomenon. Third, the retained-asset share τ_i has mass at zero (whole-bank P&As) and a continuous right tail (partial retention). This bimodality is the quantity-side analog of the mass at zero in the asset-premium distribution: the FDIC either hands the entire asset book to the acquirer at a market price, or retains a nontrivial share in receivership.

Figure 7 plots the distribution of the realized retained-asset share τ, the single quantity in the press-release data most directly tied to the structural model. Thirty percent of failures have $\tau = 0$: either the acquirer takes the whole balance sheet or the resolution is a pure payoff. Among the remaining two-thirds, τ is continuously distributed with a mean of 0.43 and a right tail extending beyond 1.0 for cases in which the FDIC retains assets valued above QBF book value after write-

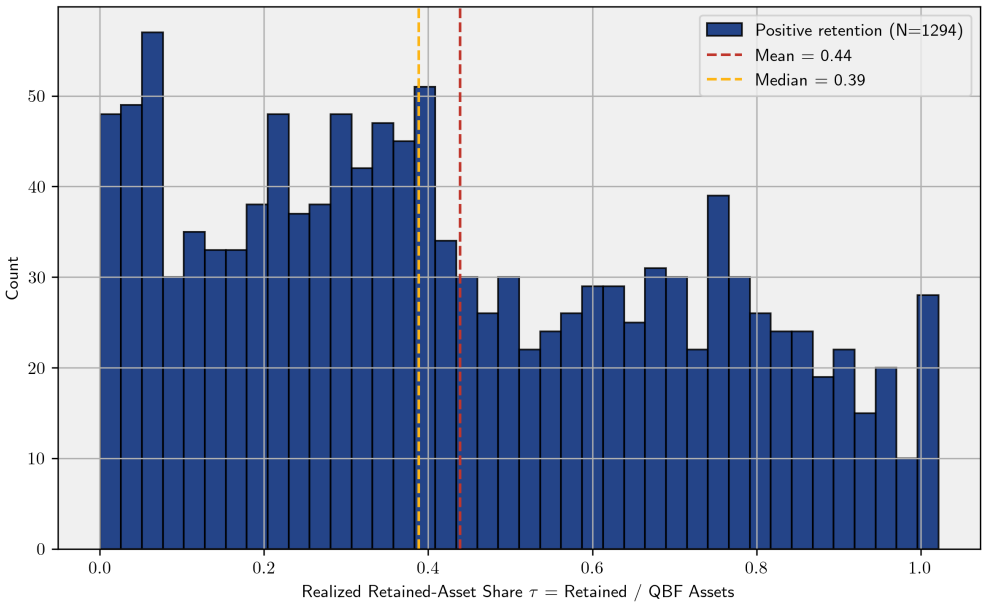


Figure 7. Distribution of the realized retained-asset share τ = Retained Assets/QBF Assets. Histogram restricted to positive retention; the fraction of failures with $\tau = 0$ is reported in the title.

downs.

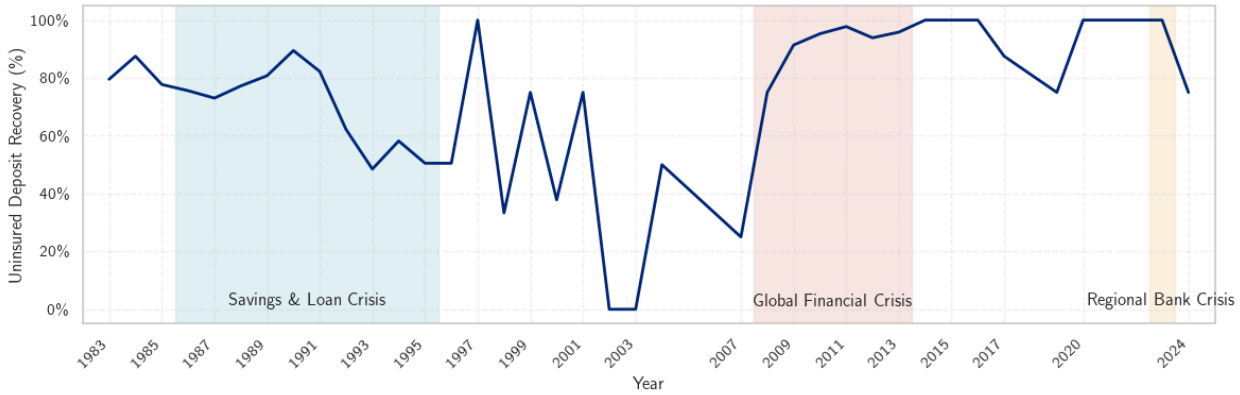


Figure 8. Uninsured deposit recovery over time. The series shows average recovery rates with shaded crisis periods: Savings & Loan crisis (blue), Global Financial Crisis (pink), and 2023 regional bank crisis (tan).

Figure 8 tracks the uninsured recovery rate over time. The unconditional mean is 91%, with 78% of failures delivering full recovery, 5% partial recovery (mean 58%), and 17% zero recovery. The time series reveals the institutional shift traced in Section II: outside major crises, recoveries are lower and more volatile, consistent with FDICIA's limits on protecting uninsured deposits absent a systemic-risk determination. In the GFC and 2023 episodes, the combination of whole-bank P&As and explicit systemic-risk exceptions pushes recoveries close to 100%. The ex-ante distribution of recoveries is a mixed incentive for uninsured depositors to run, as emphasized by Diamond and Dybvig (1983), and is empirically consistent with the stock-return evidence in Cowan and Salotti (2015) on the incidence of FDIC guarantees.

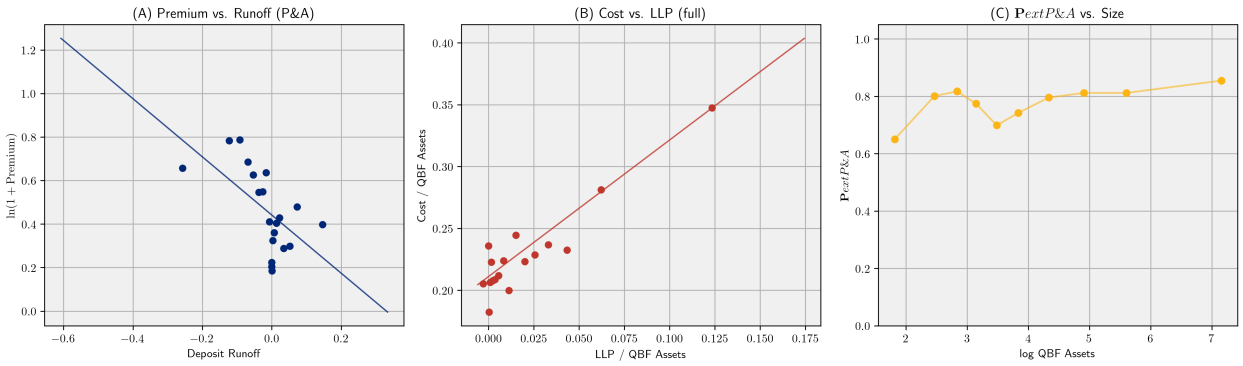


Figure 9. Motivating bivariate patterns. Panel (A): binned scatter of log deposit premium on deposit runoff, P&A sample. Panel (B): binned scatter of Cost/Assets on the LLP ratio, full sample. Panel (C): mean probability of a P&A across deciles of log QBF assets. In each panel, the line is the OLS fit (or the decile mean in Panel C) on the underlying observations.

Before turning to the structural model, Figure 9 previews the three regressions estimated in Section V with binned scatter plots of the reduced-form relationships they formalize. Panel (A) plots the log deposit premium against deposit runoff, restricted to P&A transactions: premiums decline monotonically with larger runoff, consistent with the franchise-value interpretation of the deposit premium. Panel (B) plots Cost/Assets against the LLP ratio over the full sample: the slope is strongly positive, consistent with James (1991) and Bennett and Unal (2015). Panel (C) plots the share of resolutions structured as P&As across deciles of log QBF assets: the probability of a P&A is near 0.9 in the top deciles and falls toward 0.5 for the smallest banks, reflecting the interaction of least-cost pricing with the geography of potential acquirers documented in Granja et al. (2017). These three patterns, in that order, preview Tables V, VI, and VII.

IV. Theory

This section develops a structural model of FDIC failed-bank auctions that jointly determines the deposit premium, the asset premium (or discount), and the fraction of assets retained by the FDIC. I begin with a continuous time cash flow framework for the deposit franchise, then introduce the balance-sheet objects that govern asset valuation, and embed both sides in a sealed-bid first-price scoring auction. The central analytical result is that the FDIC's least-cost mandate collapses a multidimensional bid, comprising a deposit premium, an asset purchase price, and an asset take-up share, into a scalar score, thereby preserving the tractability of standard first-price auction theory while endogenizing the composition of the resolution transaction. The model yields directly testable comparative statics linking deposit-base quality, asset toxicity, and bidder heterogeneity to observed premiums, retained asset shares, and sale probabilities.

A. *Environment*

Time t is continuous, with $t = 0$ denoting the auction date. Failed banks are indexed by $i \in \{1, \ldots, I\}$ and potential acquiring banks by $j \in \{1, \ldots, J\}$ Each failed bank i holds a stock of core deposits D_{i0} and a portfolio of assets with aggregate book value A_{i0} at the date of auction. Deposits exhibit exponential decay, at any time $t \geq 0$, the remaining deposit balance is

$$D_{it} := D_{i0} e^{-\lambda_i t}, \tag{4}$$

where $\lambda_i \in \mathbb{R}_{++}$ is the deposit-attrition hazard rate for bank i. A smaller λ_i indicates a "stickier" deposit base that runs off more slowly, whereas a larger λ_i implies rapid outflows following failure. This exponential specification is a parsimonious reduced form for the rich heterogeneity in depositor behavior documented in the literature: some depositors leave immediately upon resolution, while others remain with the acquiring institution indefinitely. The hazard rate λ_i summarizes the net effect across the depositor population.

Let r_t^m denote the market short-term interest rate and let r_{jt} denote acquirer j's marginal return on earning assets. Each acquirer also incurs an operating cost $c_j \in \mathbb{R}_+$ per dollar of deposits serviced. The deposit rate paid to depositors of failed bank i is

$$r_{it}^d := \alpha_i + \beta_i r_t^m, \tag{5}$$

where α_i is a bank-specific intercept capturing the base level of deposit compensation and $\beta_i \in [0, 1]$ measures the pass-through of market rates to deposit costs. A low β_i characterizes a deposit base that is insensitive to the prevailing rate environment, the hallmark of valuable deposit franchise. A high β_i indicates deposits that reprice rapidly with the market, eroding the funding-cost advantage that makes the franchise worth acquiring in the first place. Let $r_f \in \mathbb{R}_+$ denote the acquirer's discount rate, and define $\mu_x := \mathbf{E}[x_t]$ for any stochastic process x, so that μ_{r^m} and μ_{r_j} are the expected market rate and the expected asset return, respectively.

B. *Valuing the Deposit Franchise*

I first derive the economic value of the failed bank's deposit franchise to a potential acquirer. Because the deposit base declines over time at hazard rate λ_i[4], the franchise is a wasting asset that generates a stream of cost-saving benefits until deposits run off. At each instant, the acquirer can invest the remaining deposits in earning assets at rate r_{jt}, while paying depositors the rate r_{it}^d and bearing operating cost c_j. A single dollar of deposits therefore produces instantaneous net profit

$$\Pi_t := \left(r_{jt} - r_{it}^d - c_j \right) D_t. \tag{6}$$

I impose $\mathbf{E}[\Pi_t] > 0$, consistent with the FDIC's requirement that acquiring institutions be in

[4]The franchise is a wasting asset because the acquirer inherits a fixed, run-off deposit portfolio, no new deposits are originated to replace attrition.

sound financial condition and capable of profitably deploying the acquired deposit base.

The present value of the deposit franchise is the expected discounted flow of net spreads, accounting for exponential deposit attrition

$$V_{ij}^{D}(\beta_i, \lambda_i) := \mathbf{E}\left[\int_0^\infty e^{-r_f t}\Pi_t dt\right] = \frac{\left(\mu_{r_j} - \alpha_i - \beta_i \mu_{r^m} - c_j\right) D_{i0}}{r_f + \lambda_i}. \tag{7}$$

This closed-form expression distills the franchise value into a transparent function of three economic forces. The numerator captures the expected net interest margin per dollar of deposits: the spread between asset returns and the all-in cost of deposit funding. The denominator is an effective discount rate that combines the acquirer's time preference r_f with the deposit-attrition rate λ_i. Intuitively, deposits that decay faster are worth less because the acquirer has a shorter window over which to harvest the spread.

Two parameters are central to the empirical analysis that follows. First, a lower β_i preserves the acquirer's funding advantage even as market rates rise, directly increasing the numerator of (7). Second, a lower λ_i reduces the effective discount rate in the denominator, amplifying the present value of every dollar of expected spread. Both channels raise V_{ij}^{D} and, as I show below, feed through monotonically into equilibrium bid premiums.

C. Asset Valuation

While the deposit franchise constitutes the primary source of going-concern value in most FDIC resolutions, the asset portfolio of the failed bank also enters the acquirer's valuation. In practice, FDIC P&A transactions allow bidders to acquire some or all of the failed bank's assets alongside its deposits. Assets that are not transferred to the winning bidder are retained by the FDIC as receiver and liquidated or resolved over time. To capture this institutional feature, I introduce a minimal asset-valuation structure that generates endogenous variation in the fraction of assets transferred.

Let $r_i^{\text{FDIC}} \in (0, 1)$ denote the FDIC's expected recovery value per dollar of book assets if bank i's assets are retained and worked out in receivership. I treat r_i^{FDIC} as a reduced-form primitive that summarizes the receiver's liquidation technology, the composition of the asset portfolio (e.g., performing loans versus distressed commercial real estate), and the time and cost of the workout process. Higher r_i^{FDIC} indicates that the FDIC is relatively efficient at recovering value from the asset pool.

Let $s_{ij} \in [0, 1]$ denote the fraction of the failed bank i's total asset book that a bidder j proposes to acquire, so that $s_{ij}A_{i0}$ dollars of assets are transferred and $(1-s_{ij})A_{i0}$ are retained by the FDIC. I model each bidder j's marginal valuation for the x-th fractional unit of bank i's assets as

$$v_{ij}(x) := \bar{v}_{ij} - \gamma_i x, \qquad \gamma_i > 0. \tag{8}$$

Here $\bar{v}_{ij} \in (0, 1]$ is bidder j's marginal valuation for the first tranche of assets, and $\gamma_i > 0$ governs the rate at which marginal value declines as the bidder absorbs a larger share of the portfolio.

The declining marginal value specification in (8) is motivated by three economic forces. First, failed-bank asset portfolios are heterogeneous: performing residential mortgages, consumer auto loans, commercial real estate, and distressed or non-performing assets coexist on the same balance sheet. A rational bidder will cherry-pick the highest-quality tranches first, so marginal value falls as the bidder moves deeper into the portfolio. Second, acquiring banks face regulatory capital constraints; each additional dollar of risk-weighted assets demands incremental capital, raising the effective cost of the marginal asset. Third, absorbing a failed bank's assets requires operational capacity, loan servicing infrastructure, workout expertise, local market knowledge, that is subject to diminishing returns at scale. The slope parameter γ_i captures the combined force of these frictions. When γ_i is large, the portfolio is highly heterogeneous or the bidder's capacity constraints bind quickly, so marginal value deteriorates rapidly. When γ_i is small, the asset pool is relatively homogeneous and the bidder can absorb a large share without significant value erosion.

Integrating (8) over the take-up share s_{ij} yields the bidder's total value from acquiring fraction s_{ij} of the asset book

$$V_{ij}^{A}(s_{ij}) := A_{i0} \int_0^{s_{ij}} v_{ij}(x)dx = A_{i0} \left(\bar{v}_{ij} s_{ij} - \frac{1}{2} \gamma_i s_{ij}^2 \right). \tag{9}$$

The quadratic form in (9) has a natural interpretation; the first term $\bar{v}_{ij} s_{ij}$ is the value the bidder would obtain if all acquired assets were worth $\bar{v}_{ij}$ per dollar of book, while the second term $\frac{1}{2}\gamma_i s_{ij}^2$ is the cumulative "toxicity penalty" that reflects the declining quality of successive asset tranches. For small s_{ij}, the penalty is negligible and the bidder captures close to $\bar{v}_{ij}$ per dollar; for s_{ij} near one, the penalty can be substantial, potentially rendering the marginal asset a net liability. This curvature is the minimal device needed to generate interior solutions for asset take-up; without it, bidders would either acquire all assets or none, counterfactual to the rich variation in retained asset shares observed in the data.

D. Auction Mechanism

The FDIC resolves the majority of failed banks through P&A transactions structured as sealed-bid first-price auctions. I extend the standard single-dimensional format to accommodate joint bidding on deposits and assets, then show that the FDIC's statutory least-cost mandate reduces the multidimensional bid to a scalar score, preserving the tractability of classical auction theory.

I restrict attention to whole-bank P&A transactions in which the acquirer assumes all deposits, including uninsured deposits, and no loss-share arrangement is in place. Each bidder j submits a sealed bid for failed bank i consisting of three components:

(1) a *deposit premium rate* $p_{ij} \geq 0$, expressed per dollar of assumed deposits, so that total premium dollars are $P_{ij} = p_{ij} D_{i0}$,
(2) an *asset take-up share* $s_{ij} \in [0, 1]$, specifying the fraction of the failed bank's asset book the bidder proposes to acquire, and

(3) an *asset purchase price* $q_{ij} \in [0,1]$ per dollar of book value for the acquired assets, so that total asset payment is $Q_{ij} = q_{ij} s_{ij} A_{i0}$.

Under this bid, the acquirer pays a total of $p_{ij} D_{i0} + q_{ij} s_{ij} A_{i0}$ to the FDIC, assumes all D_{i0} of deposits as liabilities, takes ownership of $s_{ij} A_{i0}$ of assets, and the FDIC retains $(1 - s_{ij}) A_{i0}$ for disposal in receivership. The three-dimensional bid (p_{ij}, q_{ij}, s_{ij}) nests the standard deposit-only auction as the special case $s_{ij} = 0$, in which no assets are transferred and the entire premium is paid on the deposit line.

The joint bid structure reflects institutional practice: in actual FDIC resolutions, acquirers routinely negotiate over both the deposit premium and the terms on which assets are transferred. Modeling these as separate but simultaneous choices allows us to endogenize the composition of the transaction and, in particular, the fraction of assets retained by the FDIC.

FDICIA requires the FDIC to resolve each bank failure using the method that is least costly to the DIF. I formalize this mandate as a scoring rule over bids. Consider a bid (p, q, s) for failed bank i. The FDIC's cash inflows consist of the deposit premium pD_{i0} and the asset sale proceeds qsA_{i0}. Against these, the FDIC must honor the full D_{i0} of assumed deposit liabilities and expects to recover $r_i^{\text{FDIC}}(1-s)A_{i0}$ from the retained assets. The net cost to the DIF is therefore

$$C(p,q,s) := D_{i0} - \underbrace{(pD_{i0} + qsA_{i0})}_{\text{cash received from acquirer}} - \underbrace{r_i^{\text{FDIC}}(1-s)A_{i0}}_{\text{expected recovery on retained assets}}. \tag{10}$$

The first term is the deposit liability the FDIC must cover; the second is the total cash the acquirer pays; and the third is the present value the FDIC expects to realize by working out the assets it keeps. A higher bid, whether through a larger premium, a higher asset price, or a greater asset take-up share, reduces the DIF's out-of-pocket cost.

Rearranging (10) yields

$$C(p,q,s) = \underbrace{(D_{i0} - r_i^{\text{FDIC}} A_{i0})}_{\text{baseline shortfall}} - [pD_{i0} + (q - r_i^{\text{FDIC}}) sA_{i0}]$$

Assumption 1: *The failed bank is insolvent in expectation:* $D_{i0} > r_i^{FDIC} A_{i0}$.

Under Assumption 1, the baseline shortfall $D_{i0} - r_i^{\text{FDIC}} A_{i0} > 0$, is common across all bids, and reflects the fundamental insolvency of the failed bank: its deposit liabilities exceed the expected liquidation value of its assets. The FDIC's bid-selection problem therefore reduces to maximizing the second term, which I denote the *score*

$$\mathcal{S}_i(p,q,s) := pD_{i0} + (q - r_i^{\text{FDIC}})\, sA_{i0}. \tag{11}$$

This is the key tractability result of the model. Despite the three-dimensional bid space, the FDIC's least-cost mandate induces a linear scoring rule that compresses all bid components into a single scalar. I refer to $\mathcal{S}_i$ as the score—the object the FDIC evaluates when comparing bids. The score

has a clean economic interpretation. The first term, pD_{i0}, is the direct cash benefit from the deposit premium: every dollar of premium reduces DIF cost dollar-for-dollar. The second term, $(q - r_i^{\text{FDIC}})sA_{i0}$, measures the incremental benefit of asset transfer relative to the FDIC's outside option of retaining and resolving those assets itself. Transferring an additional dollar of book assets to the acquirer benefits the FDIC if and only if the acquirer's purchase price q exceeds the FDIC's own expected recovery r_i^{FDIC}. When $q < r_i^{\text{FDIC}}$, the FDIC would recover more by keeping the assets, so the asset-transfer component of the score is negative and the bidder must compensate with a higher deposit premium. This scoring structure nests the pure deposit-premium auction as the limiting case: when $s = 0$ or $q = r_i^{\text{FDIC}}$, the score collapses to pD_{i0} and the FDIC selects the highest deposit premium.

Each bidder j values the acquisition of failed bank i along two margins. On the deposit side, the bidder derives franchise value V_{ij}^D as defined in (7), scaled by a private synergy parameter $\theta_{ij} \in (0, 1]$ that captures heterogeneity in strategic fit, geographic overlap, or operational capacity. Note that V_{ij}^D carries a j subscript because it depends on acquirer-specific parameters c_j and μ_{r_j}; the franchise value of the same deposit base differs across bidders. On the asset side, the bidder obtains value $V_{ij}^A(s)$ from acquiring share s of the asset book, as defined in (9). The bidder also bears an idiosyncratic integration cost $\bar{K}_{ij} \in \mathbb{R}_+$ encompassing systems conversion, regulatory compliance, personnel retention, and other merger-related expenses.

If bidder j wins bank i's auction with bid (p, q, s), the bidder's payoff is

$$U_{ij}(p, q, s) := \theta_{ij} V_{ij}^D(\beta_i, \lambda_i) + V_{ij}^A(s) - \bar{K}_{ij} - pD_{i0} - qsA_{i0}. \tag{12}$$

The first term is the deposit franchise value scaled by synergy; the second is the gross asset value from (9); the third is the integration cost; and the final two terms are the cash payments to the FDIC. This payoff structure makes transparent the fundamental trade-off the bidder faces: paying more (higher p or q) increases the probability of winning but reduces the surplus conditional on winning.

The private-information structure is as follows. The parameters θ_{ij} and $\bar{K}_{ij}$ are private to bidder j; rival bidders and the FDIC observe only their joint distribution. The deposit-franchise parameters (β_i, λ_i) and the asset-portfolio parameters $(\bar{v}_{ij}, \gamma_i)$ are common knowledge in the sense that they characterize the failed bank's observable balance sheet, though bidders may differ in $\bar{v}_{ij}$ due to heterogeneous asset-management capabilities.

I now show that the combination of the FDIC's linear scoring rule and the bidder's payoff structure reduces the multidimensional auction to a standard first-price auction on the scalar score $\mathcal{S}_i$. This reduction is the central analytical device of the model. The first step is to express the bidder's total cash payment in terms of the score. Using (11),

$$pD_{i0} + qsA_{i0} = \mathcal{S}_i + r_i^{\text{FDIC}} sA_{i0}. \tag{13}$$

Substituting (13) into the bidder payoff (12) and collecting terms yields

$$U_{ij} = \underbrace{\left[\theta_{ij} V_{ij}^{D}(\beta_i, \lambda_i) - \bar{K}_{ij} + A_{i0}\left((\bar{v}_{ij} - r_i^{\text{FDIC}})s_{ij} - \frac{1}{2}\gamma_i s_{ij}^2\right)\right]}_{:=W_{ij}(s_{ij})} - \mathcal{S}_i. \tag{14}$$

The bracketed expression, which I denote $W_{ij}(s)$, is the bidder's *private value*: the total surplus the bidder would earn if it won the auction at a score of zero. The payoff is then the private value minus the score, which is the canonical structure of a first-price auction with private value W_{ij} and bid $\mathcal{S}_i$.

The economic content of equation (14) is as follows. The term $\theta_{ij} V_{ij}^{D} - \bar{K}_{ij}$ is the net deposit-side surplus. The asset-side contribution $A_{i0}[(\bar{v}_{ij} - r_i^{\text{FDIC}})s - 1/2\gamma_i s^2]$ measures the surplus the bidder creates by acquiring assets relative to the FDIC's own recovery. The relevant wedge is not the bidder's absolute valuation $\bar{v}_{ij}$ but the comparative advantage $\bar{v}_{ij} - r_i^{\text{FDIC}}$: the bidder creates surplus on the asset transfer only to the extent that it can extract more value from the assets than the FDIC receiver could through its own workout process. This comparative-advantage structure ensures that asset transfers occur precisely when they are socially efficient, a direct consequence of the least-cost mandate.

The reduction to a scalar auction preserves the entire apparatus of first-price auction theory. In particular, the equilibrium analysis of Milgrom and Weber (1982) applies directly: each bidder submits a score bid $\mathcal{S}_i$ that is a monotonically increasing function of their private value W_{ij}, with bid shading that depends on the distribution of rival private values. The highest-private-value bidder wins and pays a score strictly below their value, earning positive expected surplus. All comparative statics derived for standard first-price auctions carry over immediately.

E. Retained Shareof Assets

The asset take-up share s is not merely a component of the bid vector; it is a strategic choice that each bidder optimizes as part of its bid preparation. Because the private value $W_{ij}(s)$ is the bidder's "budget" for bidding aggressively in the scoring auction, each bidder selects s to maximize $W_{ij}(s)$; a higher private value translates directly into a higher equilibrium score bid, which raises the probability of winning.

Taking the first-order condition of $W_{ij}(s)$ with respect to s and solving yields the optimal asset take-up share

$$s_{ij}^{*} := \min\left\{\frac{\bar{v}_{ij} - r_i^{\text{FDIC}}}{\gamma_i}, 1\right\}, \tag{15}$$

where I impose $s_{ij}^{*} \geq 0$, a bidder cannot be forced to return assets, and $s_{ij}^{*} \leq 1$, a bidder cannot acquire more than the full portfolio. The second-order condition $\partial_s^2 W_{ij} = -A_{i0}\gamma_i < 0$ confirms that this is a global maximum.

Equation (15) encodes a sharp economic logic. The bidder acquires assets up to the point where

the marginal asset's private value $v_{ij}(s^*_{ij}) = \bar{v}_{ij} - \gamma_i s^*_{ij}$ exactly equals the FDIC's recovery value r_i^{FDIC}. Assets with private value above r_i^{FDIC} generate positive surplus when transferred, so acquiring them raises the bidder's valuation. Assets with private value below r_i^{FDIC} are worth more to the FDIC in receivership than to the bidder, so acquiring them would reduce the bidder's valuation. The optimal take-up share thus partitions the asset portfolio into a "transferable" tranche, for which private-sector management dominates FDIC technology, and a "retained" tranche, for which the reverse is true.

I define the model-implied *retained asset share* as the fraction of assets kept by the FDIC under the winning bidder's optimal take-up

$$\tau_i := 1 - s^*_{ij^*}, \tag{16}$$

where j^* denotes the auction winner. The retained share is an equilibrium object, jointly determined by the bidder's asset-management comparative advantage ($\bar{v}_{ij^*} - r_i^{\text{FDIC}}$), the heterogeneity of the asset portfolio (γ_i), and the FDIC's resolution technology (r_i^{FDIC}).

The model predicts a positive association between the retained share τ_i and the deposit premium p, operating through two reinforcing channels. The first is mechanical: using (11), the deposit premium can be expressed as $p = [\mathcal{S}_i - (q - r_i^{\text{FDIC}})sA_{i0}]/D_{i0}$, so for a given score $\mathcal{S}_i$, a lower take-up share (higher τ) shifts the bid's composition toward the premium line. The second is an equilibrium channel: when marginal assets satisfy $v_{ij}(x) < r_i^{\text{FDIC}}$, forcing the bidder to acquire them would lower W_{ij} and depress the equilibrium score bid. Allowing the FDIC to retain these negative-surplus assets removes a drag on the bidder's private value, enabling more aggressive bidding that raises both $\mathcal{S}_i$ and the deposit premium.

Crucially, both channels generate the same positive sign. This means that the model's prediction for the relationship between the retained share and the premium is robust to the joint determination of τ_i and p: regardless of whether one conditions on the score or allows the score to adjust endogenously, a higher τ_i is associated with a higher p. In the empirical analysis that follows, I construct an exogenous proxy for the retained share using pre-failure balance-sheet characteristics of the failed bank. The model provides the structural basis for this approach: cross-sectional variation in asset-portfolio quality shifts the primitives $\bar{v}_{ij}$ and γ_i, which determine s^*_{ij} and hence τ_i through (15). Because these balance-sheet characteristics are determined before the resolution auction, they are not a function of the bidding equilibrium.

F. Comparative Statics

I now derive comparative statics along three dimensions: deposit-franchise fundamentals, asset-side parameters, and the interaction between the two. These results generate the model's testable predictions.

The sensitivity of the deposit franchise value to the rate-sensitivity parameter is

$$\frac{\partial V^D_{ij}}{\partial \beta_i}(\beta_i, \lambda_i) = -\frac{\mu_{r^m} D_{i0}}{r^f + \lambda_i} < 0. \tag{17}$$

A higher β_i reduces franchise value because deposits that reprice rapidly with the market compress the net interest margin in rising-rate environments. When deposit costs track market rates one-for-one ($\beta_i \to 1$), the acquirer's funding-cost advantage vanishes and the franchise reduces to a costly liability net of operating expenses. The magnitude of the effect is scaled by the expected market rate μ_{r^m}; in high-rate environments, the penalty for rate-sensitive deposits is proportionally larger.

The sensitivity to the attrition rate is

$$\frac{\partial V_{ij}^D}{\partial \lambda_i}(\beta_i, \lambda_i) = -\frac{\mathbf{E}[\Pi_t] D_{i0}}{(r^f + \lambda_i)^2} < 0. \tag{18}$$

Faster deposit attrition reduces franchise value by shortening the horizon over which the acquirer can harvest the net spread. λ_i acts as an additive premium to the discount rate, and increasing it shrinks the present value of the entire cash-flow stream. Note that the numerator is the expected net margin, which we have assumed positive; the sign of the derivative is therefore unambiguously negative.

The cross-partial derivative reveals an important interaction:

$$\frac{\partial^2 V_{ij}^D}{\partial \beta_i \partial \lambda_i}(\beta_i, \lambda_i) = \frac{\mu_{r^m} D_{i0}}{(r^f + \lambda_i)^2} > 0. \tag{19}$$

This positive cross-derivative indicates that the marginal cost of rate sensitivity is smaller when deposits are short-lived. Intuitively, if deposits are going to run off quickly regardless, their rate sensitivity matters less because the acquirer's exposure to interest-rate risk is brief. Conversely, when deposits are very sticky, rate sensitivity becomes acutely important because the acquirer bears the compressed spread for a long horizon. Stickiness and rate insensitivity are substitutes in generating franchise value: strength in one dimension partially compensates for weakness in the other.

The optimal take-up share responds to the bidder's asset-management advantage and the portfolio's heterogeneity

$$\frac{\partial s_{ij}^*}{\partial \bar{v}_{ij}} = \frac{1}{\gamma_i} > 0, \qquad \frac{\partial s_{ij}^*}{\partial \gamma_i} = -\frac{\bar{v}_{ij} - r_i^{\text{FDIC}}}{\gamma_i^2} < 0. \tag{20}$$

A bidder with a higher baseline valuation $\bar{v}_{ij}$ acquires a larger share of the portfolio because more tranches clear the r_i^{FDIC} threshold. Conversely, greater portfolio heterogeneity (higher γ_i) causes marginal value to decline more steeply, so the bidder reaches the indifference point sooner and takes up less.

The response to the FDIC's recovery technology is

$$\frac{\partial s_{ij}^*}{\partial r_i^{\text{FDIC}}} = -\frac{1}{\gamma_i} < 0. \tag{21}$$

When the FDIC is relatively efficient at working out assets (r_i^{FDIC} high), the comparative-advantage wedge $\bar{v}_{ij} - r_i^{\text{FDIC}}$ narrows, and the bidder optimally takes up a smaller share. In the limit $r_i^{\text{FDIC}} \to \bar{v}_{ij}$,

the bidder has no comparative advantage and the entire asset book is retained. This result has a clear policy implication: investments in the FDIC's receivership infrastructure (e.g., loan-workout expertise, technology for managing distressed portfolios) endogenously shift the equilibrium toward greater retention, which may or may not reduce total resolution cost depending on the deposit-side response.

Substituting the interior optimal take-up s_{ij}^* from (15) into the private value W_{ij} yields the maximized private value

$$W_{ij}(s_{ij}^*) = \theta_{ij} V_{ij}^D(\beta_i, \lambda_i) - \bar{K}_{ij} + \frac{A_{i0}(\bar{v}_{ij} - r_i^{\text{FDIC}})^2}{2\gamma_i}. \tag{22}$$

The three terms correspond, respectively, to the net deposit-side surplus, the integration cost, and the asset-side option value, the surplus the bidder extracts from acquiring those asset tranches for which it has a comparative advantage over the FDIC. The asset-side option value is always non-negative and is quadratic in the comparative-advantage wedge $\bar{v}_{ij} - r_i^{\text{FDIC}}$, reflecting the fact that a larger wedge both increases the surplus on each infra-marginal tranche and expands the set of tranches worth acquiring.

Equation (22) reveals that the deposit and asset channels are additively separable in the private value. This separability has two important implications. First, improvements in deposit-franchise quality increase the private value without altering the optimal asset take-up, because s_{ij}^* depends only on asset-side parameters. Second, improvements in asset-side comparative advantage raise the private value without altering the deposit-side valuation. The two channels reinforce each other in driving up equilibrium score bids and premiums, but they operate through independent mechanisms. This additive structure is a consequence of the linear scoring rule and the separability of the bidder's payoff across deposit and asset margins.

G. Probability of P&A

Not every failed bank is resolved through a P&A transaction. When no bidder submits a bid that clears the FDIC's least-cost threshold, the FDIC instead executes a deposit payoff (PO): insured depositors are paid directly from the DIF, uninsured depositors receive a receivership claim, and all assets are liquidated by the FDIC as receiver. The model generates predictions about which failures result in a P&A sale and which do not.

A P&A transaction is least-cost relative to a payoff if and only if the winning bid's DIF cost from (10) is below the cost of the payoff alternative. Let C_i^{payoff} denote the DIF cost under a payoff resolution. A sale occurs if and only if

$$\max_j \mathcal{S}_i(p_{ij}, q_{ij}, s_{ij}) \geq \underbrace{(D_{i0} - r_i^{\text{FDIC}} A_{i0}) - C_i^{\text{payoff}}}_{:=\underline{\mathcal{S}}_i}, \tag{23}$$

where $\underline{\mathcal{S}}_i$ is the minimum score required for the P&A to be cost-effective. When the best available score falls below this threshold, no sale occurs.

In equilibrium, each bidder's score bid is a monotonically increasing function of their maximized private value $W_{ij}(s_{ij}^*)$ from (22). The probability that at least one bidder clears the threshold $\underline{\mathcal{S}}_i$ therefore depends on the distribution of private values across potential acquirers.

Several predictions follow directly. First, failures with higher deposit-franchise value are more likely to attract bids that clear the threshold, because V_{ij}^D enters positively in every bidder's private value. Attractive deposit franchises are, in a precise sense, easier to sell. Second, failures with more toxic asset portfolios present a subtler prediction. The direct effect of asset toxicity is to reduce the asset-side option value in (22), which lowers private values and reduces sale probability. However, the endogenous retention mechanism partially offsets this: by allowing the FDIC to retain the most toxic tranches, the model removes the negative-surplus assets from the bidder's calculus, sustaining a positive private value even when the overall portfolio is impaired. The net effect on sale probability depends on the magnitude of the comparative-advantage wedge $\bar{v}_{ij} - r_i^{\text{FDIC}}$ relative to γ_i. Third, the number of bidders J affects sale probability through the standard order-statistic channel: the maximum score bid among J draws from the private-value distribution is stochastically increasing in J. Greater competition raises the probability that at least one bidder clears the threshold, consistent with the well-established empirical finding that FDIC failures attracting more interested parties are more likely to be resolved through P&A transactions.

Among failures that result in a P&A sale, the model generates a positive conditional association between the retained share τ_i and the deposit premium p. Failures that clear the sale threshold despite high asset toxicity must have particularly valuable deposit franchises or particularly well-matched acquirers. Conditional on sale, the high-τ observations are therefore selected to have strong deposit-side fundamentals, which manifest as high premiums. This selection effect reinforces the mechanical and equilibrium channels identified in Section IV.E.

H. Testable Predictions

The model generates six testable predictions that organize the empirical analysis in the following section.

The first two predictions concern the deposit franchise. The winning deposit premium p is decreasing in the deposit-rate sensitivity β_i, because rate-insensitive deposits preserve the acquirer's funding-cost advantage in rising-rate environments, raising franchise value and equilibrium bids. The premium is also decreasing in the deposit-attrition rate λ_i, because stickier deposits extend the horizon over which the acquirer earns the net spread, amplifying the present value of the franchise through the effective discount rate $r_f + \lambda_i$.

The next three predictions concern the asset side and its interaction with the premium. The retained asset share τ_i is increasing in asset-portfolio heterogeneity γ_i and decreasing in the bidder's asset-management advantage $\bar{v}_{ij}$: the FDIC retains a larger fraction of assets when the portfolio is more heterogeneous or when bidders have less comparative advantage in managing distressed assets. Conditional on a P&A sale, the deposit premium p is positively associated with the retained share τ_i, a prediction that holds through both the mechanical and equilibrium channels identified in Section

IV.E and that the model's structure signs unambiguously. The retained share is also increasing in the FDIC's expected recovery value r_i^{FDIC}, because a more efficient FDIC receiver raises the threshold for asset transfer, leading to greater equilibrium retention.

The final prediction concerns the extensive margin of resolution. The probability of a P&A transaction, relative to a deposit payoff, is increasing in the deposit-franchise value V_{ij}^D, the asset-side comparative-advantage wedge $\bar{v}_{ij} - r_i^{\text{FDIC}}$, and the number of potential bidders J. Each operates through the private value $W_{ij}(s_{ij}^*)$ and the order-statistic structure of the auction: more valuable franchises, larger asset-side surpluses, and greater competition all raise the probability that the maximum score bid exceeds the least-cost threshold $\underline{\mathcal{S}}_i$.

The model is intentionally parsimonious. It abstracts from several features of actual FDIC resolutions: loss-share agreements, partial deposit assumptions, dynamic renegotiation between the FDIC and potential acquirers, and the possibility that the FDIC's recovery value r_i^{FDIC} is itself endogenous to the resolution method. These omissions are deliberate: the goal is to isolate the first-order interaction between deposit-franchise quality and asset-side composition in a framework that is transparent enough to generate sharp, sign-identified predictions. The additive separability of the deposit and asset channels in the maximized private value (22) is the key structural result, it implies that deposit-side and asset-side fundamentals can be tested independently in the cross-section, and that their contributions to the premium can be decomposed without requiring a fully specified joint distribution of bidder types. The empirical specifications in the next section map each prediction to an observable regressor, exploiting pre-failure balance-sheet variation to identify the model's comparative statics.

V. Empirics

The theoretical framework of Section IV maps six comparative statics onto three estimating tests: a deposit pricing channel, an FDIC cost channel, and a selection channel. The analysis centers on three constructed variables. Deposit runoff proxies for the attrition parameter λ_i and is defined in (1). Deposit beta measures the interest-rate sensitivity of deposit costs and is matched from Drechsler et al. (2021). The toxic ratio proxies for the retained asset share τ_i and equals the fraction of total assets held by the FDIC receivership rather than transferred to the acquirer.

Two identification threats require explicit attention before turning to the results. The first is simultaneity in the toxic ratio: because the premium and the retained asset share are jointly determined at the auction, OLS estimates of the toxic ratio coefficient in the pricing regression are biased toward zero. I address this by instrumenting the realized toxic ratio with the pre-failure loan loss provision ratio, defined in (2).[5] The first-stage coefficient on the provision ratio is 0.897 with a standard error of 0.231, yielding an F-statistic of 15.0, which satisfies the Staiger and Stock (1997) threshold for instrument relevance. The first stage reflects a mechanical relationship: banks with

[5]The loan loss provision ratio is the ratio of total loan loss provisions to assets in the quarter before failure. It is determined on the balance sheet before any bidding occurs and is therefore unaffected by the auction outcome.

larger pre-failure provisioning needs generate a larger pool of impaired assets that acquirers decline to purchase, so more of the asset book remains with the FDIC. The full first-stage regression is reported in Table X of Appendix D.B.

The second identification threat is sample selection on P&A transactions: a deposit premium is observed only when a sale occurs, so the pricing regression is estimated on a subsample that is not a random draw from the failure universe. I address this with the two-step correction of Heckman, Ichimura, and Todd (1997), using the log number of concurrent failures in the same calendar quarter as the exclusion restriction. The argument for exclusion follows James (1991) and Granja et al. (2017): resolution bandwidth and acquirer attention are finite, so aggregate quarterly congestion reduces the probability of a timely P&A bid without affecting the conditional price that a successful acquirer pays for the specific target. The Heckman correction leaves the pricing coefficients essentially unchanged; the selection-corrected estimates and full diagnostics are reported in Appendix D.B.

The headline result across all three tests is that the toxic ratio is the dominant quantitative lever. Deposit runoff and deposit beta enter with the signs predicted by the model but with attenuated statistical significance.

A. Data and Sample

The sample spans 1,860 failure observations from 1983 through 2024 across four distinct regulatory epochs. Sample coverage varies systematically across the three tests. The full failure universe supports the cost and selection regressions, with between 1,784 and 1,860 observations depending on variable availability. The P&A subsample, restricted to transactions for which an explicit deposit premium is observed, provides approximately 1,488 observations for the pricing regressions. The deposit-beta-matched subsample further restricts to banks for which Drechsler et al. (2021) supply a matched rate-sensitivity estimate, yielding 373 observations. An uninsured-share subsample of 133 observations, concentrated in the pre-1990 period, isolates the compositional effects of liability structure.

Year fixed effects appear in every specification and robust standard errors are used throughout. State fixed effects are omitted on power grounds given the sample size.[6]

B. Deposit Pricing Channel

The pricing regression implements the comparative statics of the deposit franchise value in equations (17) and (18). The estimating equation is

$$\ln\left(1+\frac{\text{Premium}_{it}}{\text{Deposits}_{it}}\right)=\alpha+\gamma_1\text{Deposit Runoff}_i+\gamma_2\text{Deposit Beta}_i+\Gamma X_i+\delta_t+\varepsilon_{it}, \tag{24}$$

[6]State fixed effects would ideally absorb idiosyncratic regional shocks such as the collapse of oil prices in the 1980s, which disproportionately affected community banks in Texas and Oklahoma. The sample does not have sufficient within-state variation to support them.

where X_i includes the loan loss provision ratio and log asset size and δ_t are year fixed effects. The model predicts $\gamma_1 < 0$ and $\gamma_2 < 0$. Table V reports the results.

	Deposit Premium: $\ln(1+p)$				
	(1)	(2)	(3)	(4)	(5)
Deposit Runoff	0.00763	0.01628	0.01589	0.06707***	0.01623
	(0.00518)	(0.01177)	(0.01176)	(0.02066)	(0.01182)
Uninsured Share				0.06238***	
				(0.01696)	
Deposit Beta		-0.01468			-0.02052
		(0.00900)			(0.01408)
Expense Beta			-0.01221		0.00723
			(0.00931)		(0.01413)
LLP Ratio	-0.00508	0.04591	0.04559	0.07518	0.04592
	(0.01331)	(0.04634)	(0.04712)	(0.09179)	(0.04608)
Log Asset Size	0.00033	-0.00088	-0.00088	-0.00029	-0.00091
	(0.00035)	(0.00056)	(0.00056)	(0.00143)	(0.00056)
Year FE	X	X	X	X	X
N	1488	359	359	133	359
R^2	0.30	0.45	0.45	0.46	0.45

Table V: Deposit Side: Pricing Channel. OLS regressions of the log deposit premium on deposit-franchise characteristics and the exogenous LLP ratio proxy. Sample restricted to P&A transactions. All specifications include year fixed effects. Robust standard errors in parentheses. Significance levels: $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

The primary finding of the pricing regressions is the narrow dispersion of the dependent variable. The sample standard deviation of $\ln(1+p)$ on the P&A subsample is approximately 0.015 log-points, and year fixed effects absorb most of the cross-sectional variation, yielding $R^2 = 0.30$ in column (1). This compression reflects the competitive structure of the auction: the winning bid is tightly concentrated around a year-specific median, leaving little residual variation for deposit fundamentals to explain.

In column (1), deposit runoff enters with a point estimate of +0.0076 that is not statistically distinguishable from zero. Columns (2) and (3) introduce deposit beta and expense beta on the matched subsample of 359 observations. Deposit beta enters at -0.015 and expense beta at -0.012, both carrying the sign predicted by equation (17): higher rate sensitivity reduces the premium an acquirer is willing to pay. Neither achieves conventional significance. The most likely explanation is classical measurement error in rate-sensitivity estimates for small failing banks, which typically have limited exposure to the branch-level deposit pricing data underlying the Drechsler et al. (2021) betas.

Column (4), which introduces the uninsured-share subsample, generates the sharpest deposit-side pricing result in the table. Deposit runoff enters at +0.067 and uninsured share enters at

+0.062, both significant at the one-percent level. The sign on deposit runoff in this column requires interpretation. When uninsured share is controlled, residual runoff conditional on a given uninsured fraction reflects sophisticated-depositor sorting rather than generalized panic. Acquirers value the insured core franchise; runoff concentrated among uninsured depositors who bear losses in any resolution is a qualitatively different event from runoff among insured depositors whose claims the acquirer assumes. The positive coefficient is therefore consistent with the insured franchise being more valuable precisely in banks where the uninsured depositor base has already self-selected out.

The loan loss provision ratio enters with a negative point estimate in column (1). This sign reflects the selection and bargaining channel: banks that require a larger loss-sharing pool to attract any bidder are precisely the banks whose residual franchise commands a lower premium. The structural positive effect of asset cleanliness on the premium emerges once simultaneity is purged by the instrumental-variable strategy described above. The second-stage IV coefficient on the instrumented toxic ratio is larger in magnitude than the corresponding OLS coefficient, consistent with simultaneity bias attenuating OLS estimates toward zero.

C. FDIC Cost Channel

The FDIC resolution cost is the mirror image of the deposit premium. Every dollar the acquirer does not pay, the insurance fund absorbs. The estimating equation is

$$\frac{\text{Cost to FDIC}_{it}}{\text{Assets}_{it}} = \alpha + \psi_1 \text{Deposit Runoff}_i + \psi_2 \text{Deposit Beta}_i + \Psi X_i + \delta_t + \varepsilon_{it}, \tag{25}$$

where X_i includes the loan loss provision ratio and log asset size. The model predicts $\psi_1 > 0$ and $\psi_3 > 0$, with all signs opposite to those in the pricing channel. Table VI reports the results.

The cost regressions deliver the sharpest results in the paper. In column (1) all three primary regressors enter with the predicted signs at one-percent significance. Deposit runoff carries a coefficient of +0.168: a one-standard-deviation increase in runoff raises the cost-to-assets ratio by roughly 4.9 percentage points relative to a sample median of approximately 22 percent.[7] The loan loss provision ratio carries a coefficient of +0.859: a one-standard-deviation increase in the provision pool raises the cost-to-assets ratio by approximately 20 percentage points, making it the dominant quantitative driver across all specifications. Log asset size enters at −0.016: larger banks are resolved at lower proportional cost, consistent with scale economies in resolution and with the size effects documented in the selection regression below.

These results confirm the theoretical trade-off at the core of the paper. The same provision ratio that enters negatively and insignificantly in the deposit pricing regressions, enters positively and significantly in the cost regressions at a magnitude of roughly +0.86. The same variable, measured from the same balance sheet, takes opposite signs on the two sides of the auction ledger in exactly the pattern that the model's scoring mechanism in equation (11) predicts. The FDIC's asset-retention decision simultaneously raises the acquirer's willingness to pay and raises the insurance

[7] The one-standard-deviation increase corresponds to approximately 0.29 log-points of deposit runoff.

	Cost/Assets			
	(1)	(2)	(3)	(4)
Deposit Runoff	0.16784***	0.22858*	0.22833*	0.10202
	(0.04336)	(0.12515)	(0.12540)	(0.08042)
Uninsured Share				-0.00034***
				(0.00008)
Deposit Beta		0.05810		
		(0.08943)		
Expense Beta			0.06870	
			(0.09640)	
LLP Ratio	0.85942***	0.62922**	0.62754**	0.81910***
	(0.10276)	(0.29927)	(0.29891)	(0.23310)
Log Asset Size	-0.01557***	-0.02609***	-0.02636***	-0.00487
	(0.00267)	(0.00541)	(0.00551)	(0.00604)
Year FE	X	X	X	X
N	1784	372	372	402
R^2	0.16	0.21	0.21	0.26

Table VI: OLS regressions of the cost-to-assets ratio on deposit-franchise characteristics and the exogenous LLP ratio. All specifications include year fixed effects. Robust standard errors in parentheses. Significance levels: $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

fund's ultimate loss.

Column (4) introduces the uninsured-share subsample. Uninsured share enters at -0.00034 with high significance. This is a mechanical accounting channel distinct from the franchise-value channel in the pricing regressions: in a liquidation or partial assumption, uninsured depositors absorb losses before the insurance fund, directly reducing the DIF's out-of-pocket cost. Deposit beta enters in columns (2) and (3) with a positive point estimate on the cost side, directionally consistent with the model, but neither coefficient achieves conventional significance in the smaller matched subsample.

D. P&A Selection Channel

The extensive margin of resolution, whether the FDIC finds a P&A acquirer at all, is governed by equation (23) of the model. The probability that the highest private value clears the least-cost threshold $\underline{\mathcal{S}}_i$ is increasing in the deposit franchise value V_{ij}^D, the asset-side comparative-advantage wedge, and the number of bidders. The estimating equation is

$$\mathbf{P}\{\text{Resolution Type}_i = \text{P\&A}\} = \alpha + \phi_1 \text{Deposit Runoff}_i + \phi_2 \text{Core Deposit Share}_i + \phi_3 \text{LLP Ratio}_i + \Phi X_i + \delta_t + \varepsilon_{it}, \tag{26}$$

where the dependent variable equals one if the failure was resolved via P&A or Insured Deposit Transfer and zero for deposit payoff or DINB resolutions. Table VII reports linear probability model and logit estimates.

	Linear Probability Model			Logit
	(1)	(2)	(3)	(4) ME
Deposit Runoff	-0.28891***	-0.11663	0.07859	-0.25095***
	(0.10179)	(0.10990)	(0.22185)	(0.08850)
Core Deposit Share			0.00048***	
			(0.00011)	
LLP Ratio	0.24152	0.30955	-1.58490**	0.21558
	(0.32063)	(0.33020)	(0.62798)	(0.31572)
Log Asset Size	0.03587***	0.00958	0.12576***	0.04112***
	(0.00544)	(0.00745)	(0.01493)	(0.00686)
Year FE		X		
N	1784	1784	402	1784
R^2 / Pseudo R^2	0.03	0.07	0.15	0.04

Table VII: Columns (1)–(3) report Linear Probability Model estimates. Column (4) reports logit marginal effects evaluated at the mean. Robust standard errors in parentheses. Significance: $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Column (1) of Table VII presents the baseline linear probability model without year fixed effects. Deposit runoff enters at -0.289 and log asset size enters at $+0.036$, both significant at the one-percent level. Faster pre-failure deposit attrition materially reduces the probability that any bidder submits a P&A-viable offer, consistent with the franchise-value channel in the model.

Column (2) adds year fixed effects. The runoff coefficient collapses to -0.117 and loses significance. The runoff-to-P&A relationship is primarily a cross-episode phenomenon driven by aggregate deposit flight during crisis years. Year effects absorb the macroeconomic component of deposit flight common to all failing banks in a given year, leaving only idiosyncratic runoff variation whose marginal explanatory power for resolution type is limited.

Column (3) introduces core deposit share, defined as insured deposits as a fraction of total deposits, on the uninsured-share subsample. Core deposit share enters at $+0.00048$ with high significance, confirming that the stickiness and insurance-coverage properties of the liability base

emphasized by Drechsler et al. (2021) translate directly into resolution marketability. The loan loss provision ratio reverses sign to -1.585, significant at five percent. Without core deposit share as a control, the provision ratio proxies jointly for asset-quality deterioration and for the FDIC's willingness to engineer a deal through loss-sharing arrangements. Once the franchise-quality component is absorbed by core deposit share, the provision ratio isolates the pure bad-asset component, which reduces P&A marketability. Column (4) reports logit marginal effects evaluated at the sample mean and confirms the linear probability model sign on runoff at -0.251, significant at the one-percent level.

Log asset size is the single most robust selection predictor across all columns. Larger banks are disproportionately more likely to attract P&A acquirers, consistent with scale economies in due diligence and with the historical record that the FDIC has expended substantial effort to avoid liquidating large institutions during systemic episodes.

E. Heckman Corrected Deposit Pricing Channel

As described in the opening section, the pricing regression is estimated on the P&A subsample and the Heckman correction assesses whether this conditioning generates material bias in the pricing coefficients. Table VIII presents the key results. Column (1) reproduces the baseline OLS pricing regression for comparison; Columns (2) through (5) report the second-stage estimates of the Heckman et al. (1997) correction under alternative specifications, with case-bootstrap standard errors.

The columns differ in the exclusion restriction used to identify the first-stage probit and, in Column (5), the estimation sample. Column (2) (Congestion IV) uses the log number of concurrent failures in the same calendar quarter as the sole exclusion variable. The identification argument follows James (1991) and Granja et al. (2017): when many banks are failing simultaneously, the FDIC's finite resolution bandwidth and potential acquirers' finite attention reduce the probability that a specific bank attracts a timely P&A bid, but conditional on a sale occurring, the transaction price for that specific franchise is unaffected by how many other auctions are running in the same period. Column (3) (Cycle IV) replaces the congestion measure with the log number of days elapsed since the nearest macroeconomic crisis trough, a continuous measure of where any given failure falls in the credit cycle. The cycle variable captures a different dimension of resolution-market conditions: during troughs, acquirer balance sheets are most impaired and bidder participation is thinnest, reducing the probability of a P&A outcome for reasons unrelated to the target's own pricing fundamentals. Column (4) includes both exclusion variables simultaneously in an overidentified specification; the pricing coefficients serve as a robustness check on whether the two instruments are capturing the same or independent variation in selection. Column (5) restricts to the deposit-beta matched subsample of 359 observations and adds deposit beta as both a pricing covariate and a selection covariate in the first stage.

The pricing coefficients are essentially unchanged by the correction. The deposit runoff coefficient moves from 0.0076 in OLS to 0.0065 in the primary corrected specification (column 2); the

	$\ln(1+p)$ on the P&A Subsample				
	(1) OLS	(2) Congestion IV	(3) Cycle IV	(4) Over-Id	(5) β
Deposit Runoff	0.00763	0.00648	0.00819	0.00685	0.01676
	(0.00518)	(0.00757)	(0.00871)	(0.00714)	(0.01463)
LLP Ratio	-0.00508	-0.00249	-0.00622	-0.00335	0.05407
	(0.01331)	(0.01801)	(0.02463)	(0.01500)	(0.05832)
Log Asset Size	0.00033	0.00040	0.00029	0.00038	-0.00087
	(0.00035)	(0.00046)	(0.00055)	(0.00041)	(0.00069)
Deposit Beta					-0.01387
					(0.00941)
Inverse Mills Ratio		0.00628	-0.00303	0.00422	0.03048
		(0.02359)	(0.04005)	(0.01615)	(0.05428)
Implied $\hat{\rho}$		0.419	-0.207	0.286	1.794
Implied $\hat{\sigma}_\varepsilon$		0.015	0.015	0.015	0.017
Year FE	X	X	X	X	X
N	1488	1488	1488	1488	359
R^2	0.297	0.297	0.297	0.297	0.453

Table VIII: OLS and Heckman two-step second-stage estimates of the log deposit premium on deposit-franchise controls, asset-quality proxy, and the inverse Mills ratio. Column (1) reports OLS on the P&A subsample without selection correction. Columns (2)–(5) report the second stage of a Heckman et al. (1997) two-step estimator, including the inverse Mills ratio computed from a first-stage probit of P&A selection estimated on the universe of failures. Standard errors in parentheses are heteroskedasticity-consistent in column (1) and case-bootstrapped (200 replications) in columns (2)–(5) to account for first-stage generated-regressor variability. Significance $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

provision ratio coefficient moves from -0.0051 to -0.0025; log asset size is stable at approximately 0.0004. The inverse Mills ratio coefficient is 0.0063 with a bootstrap standard error of 0.0236, failing to reject the null of no selection bias at any conventional level. The key robustness check is the stability of the pricing coefficients between columns (2) and (3): despite the two exclusion variables operating through different economic channels—within-year congestion versus cross-year cycle position—the pricing estimates barely move, which strengthens the case that the conditional pricing equation is genuinely orthogonal to resolution-market conditions. The implied correlation between the pricing and selection residuals is approximately 0.42 in the primary specification, but the corresponding bias in the pricing coefficients is bounded by a quantity of order 0.008, which is comparable to the standard errors on the pricing coefficients themselves.[8]

The economic reason the correction is small is the narrow dispersion of the dependent variable. The sample standard deviation of the log deposit premium is approximately 0.015 log-points. Even a moderate selection correlation translates into a negligible absolute bias when the outcome variable

[8]The bound is computed as $|\hat{\rho}|\hat{\sigma}_\varepsilon\bar{\lambda} \approx 0.42 \times 0.015 \times 1.3 \approx 0.008$, where $\bar{\lambda}$ is the average inverse Mills ratio on the P&A subsample.

has little total variance to begin with. This narrow dispersion is direct evidence that the FDIC auction delivers a competitive outcome in which the winning bid extracts most of the franchise surplus. The Heckman correction and the instrumental-variable strategy are complementary: the former addresses selection into the P&A sample, the latter addresses the simultaneity of the premium and the retained asset share at the auction. Full diagnostics, including the first-stage probit estimates and the decomposition of the selection correlation, are reported in Appendix D.B.

F. Summary

The three tests converge on a coherent description of how asset retention mediates the FDIC's resolution problem. The toxic ratio is the mechanical lever. Retaining distressed assets in the receivership raises the deposit premium that acquirers are willing to pay, confirmed by the 2SLS estimate in the deposit pricing channel. It raises the probability that a buyer is found at all, through the loan loss provision and core deposit share channels in the P&A selection channel. And it raises the fiscal cost borne by the Deposit Insurance Fund, with a loan loss provision coefficient of +0.859 in the FDIC cost channel. All three channels carry the predicted sign at economically large magnitudes. The Heckman correction confirms that the pricing results are not an artefact of the P&A-only sample restriction.

VI. Conclusion

This paper opens the black box of bank resolution by introducing a novel dataset of transaction terms from over forty years of FDIC press releases, transforming qualitative disclosures into a structured economic analysis of 1,860 failures. The central finding is that the apparent standardization of the modern resolution process masks a fundamental shift in regulatory technology. The transition from the chaotic liquidations of the 1980s to the monolithic dominance of the Purchase and Assumption transaction is not merely a function of industry consolidation or improved bidder appetite. It is the product of financial engineering. By retaining the most impaired assets in receivership regulators systematically cleanse distressed institutions of their troubled assets and convert them into marketable franchises. A structural model that encompasses both the liability and asset sides of the auction formalizes this mechanism, showing that the efficient partition of the asset portfolio between private and public management is a direct implication of the FDIC's least-cost mandate. The efficiency gain, however, is not free, the insurance fund absorbs the credit risk that private acquirers decline to carry.

The empirical record confirms this trade-off with consistency across all three channels of the analysis. Asset quality, as measured by the pre-failure loan-loss provision ratio, is the dominant quantitative determinant of resolution outcomes. It raises the cost borne by the Deposit Insurance Fund, depresses bid premiums unless offset by asset stripping, and reduces the probability of attracting a P&A acquirer at all. An instrumental-variable strategy that exploits the predetermined nature of the provision ratio relative to the auction confirms that the positive relationship between

asset retention and the deposit premium is causal rather than spurious. Deposit franchise quality matters in the direction the model predicts, but its quantitative influence is secondary to the asset-quality channel. The selection margin is dominated by bank size, with larger institutions resolving via purchase and assumption at rates disproportionate to their franchise fundamentals.

These findings carry direct implications for the design of the regulatory safety net. The reliance on asset retention reveals that the least-cost requirement operates with significant elasticity, effectively trading immediate marketability for retained credit risk. The fiscal cost of stripping distressed assets from a failing bank is not borne at closure; it accumulates silently on the FDIC's receivership balance sheet, socialized over the workout horizon. To the extent that market participants anticipate that the FDIC will engineer a sale for any sufficiently impaired institution by stripping its bad assets, this practice creates a distress-subsidy expectation that weakens market discipline and encourages risk-taking. The dominance of size in the selection regression adds a further dimension to this moral hazard: large institutions enjoy implicit P&A protection at rates their deposit-franchise quality alone would not justify, a quantitative manifestation of the Too-Big-to-Fail doctrine.

Mitigating these distortions requires intervention earlier in the failure process. Policymakers should consider integrating high-frequency metrics of deposit stability directly into prompt corrective action triggers. The evidence here shows that accelerated deposit attrition prior to closure destroys franchise value in a manner that is predictable *ex ante* and costly *ex post*. An eroded deposit base forces the FDIC to offer larger asset concessions to attract any bid at all. Intervening before the liability base is hollowed out would maximize auction premiums and reduce the necessity of offering expensive asset guarantees. The cyclical pattern of uninsured depositor protection compounds this problem. Systemic episodes yield near-total recovery through whole-bank P&A transactions, creating an expectation of *de facto* insurance coverage that persists into normal times and minimizes the market discipline that uninsured depositors are designed to provide. As the banking system enters a new era of digital deposit flight, where information spreads instantly and withdrawals can be executed in seconds, recalibrating these intervention thresholds is not a peripheral concern. It is the central challenge of resolution policy.

Appendix A. Data Sources

Variable	Description	Source
cert	FDIC certificate number	BankFind Suite
chclass1	FDIC bank charter class (N, NM, SM, SB)	BankFind Suite
cityst	City, state of the failed bank	BankFind Suite
cost	Cost of the resolution to the FDIC	BankFind Suite
faildate	Date the bank failed	BankFind Suite
fin	Financial institution number	BankFind Suite
id	FDIC specific ID number	BankFind Suite
name	Name of the failed bank	BankFind Suite
qbfasset	Total assets the quarter before failure	BankFind Suite
qbfdep	Total deposits the quarter before failure	BankFind Suite
restype1	The resolution type used by the FDIC (PA, IDT, PO, PI, DINB)	BankFind Suite
savr	Insurance fund	BankFind Suite
aq_name	Name of the aquiring bank	FDIC Press Releases
aq_city	City of the aquiring bank	FDIC Press Releases
aq_st	State of the aquiring bank	FDIC Press Releases
recovery	Trinary indicator for whether the deposit recovery to uninsured depositors was (F)ull, (P)artial, or (N)one	FDIC Press Releases
recovery_pct	The percentage recovery of uninsured depositors	FDIC Press Releases
fasset	Total assets of the failed bank at the time of failure	FDIC Press Releases
fdep	Total deposits of the failed bank at the time of failure	FDIC Press Releases
tot_accts	Total number of deposit accounts at the failed bank	FDIC Press Releases
amount_over	Total dollar amount of deposits over the deposit insurance limit at the failed bank	FDIC Press Releases
accts_over	Total number of accounts over the deposit insurance limit	FDIC Press Releases
prem_dep	Premium paid to aquire the deposits of the failed bank's deposits	FDIC Press Releases
prem_dep_pct	Premium paid to aquire the deposits of the failed bank's deposits as a percentage of total deposits	FDIC Press Releases
purchased_assets	Assets of the failed bank purchased by the aquiring bank	FDIC Press Releases
asset_prem	Premium paid to aquire the deposits of the failed bank's assets	FDIC Press Releases
asset_prem_pct	Premium paid to aquire the deposits of the failed bank's assets as a percentage of total assets	FDIC Press Releases
fdic_advance	FDIC laon given to the aquiring institution to facilitate the transaction	FDIC Press Releases
retianed_assets	Assets of the failed bank retained by the FDIC	FDIC Press Releases
loan_loss	Total loan loss provision to aquire the fail bank	FDIC Press Releases
loan_loss_fdic_share	The FDIC's share of the total loan loss provision	FDIC Press Releases

Table IX: This table provides the variable names, description, and source for the novel set of bank auction characterisitcs collected and constrcuted as part of this project. The BankFind Suite: Bank Failures & Assistance Data is hosted by the FDIC.

Appendix B. FDIC Failed Bank Press Release

FDIC
FEDERAL DEPOSIT INSURANCE CORPORATION

NEWS RELEASE

FOR IMMEDIATE RELEASE
PR-33-93 (4-8-93)

Media Contact:
Andrew Porterfield (202) 898-6593

FDIC APPROVES ASSUMPTION OF INSURED DEPOSITS OF
PREMIER BANK, NORTHRIDGE, CALIFORNIA

The Board of Directors of the Federal Deposit Insurance Corporation has approved the assumption of the insured deposits of Premier Bank, Northridge, California, by Trans-World Bank, Sherman Oaks, California.

The failed bank's three offices will reopen on Friday, April 9, 1993, as branches of Trans-World, and its depositors automatically will become depositors of the assuming bank.

Premier Bank, with total assets of $74.3 million, was closed on Thursday, April 8, 1993, by James E. Gilleran, California Banking Superintendent, and the FDIC was named receiver.

Trans-World will assume about $62.8 million in about 4,600 deposit accounts. At the time the bank closed, about $6.9 million in 71 accounts exceeded the federal insurance limit of $100,000 and will not be assumed by Trans-World Bank.

The Board of Directors also voted to make a prompt advance payment to uninsured depositors equal to 57 percent of the uninsured claims. Arrangements for picking up the advance payment checks can be made at any of the failed bank's offices beginning Friday, April 9, 1993.

The assuming bank will pay a premium of $701,000 for the right to receive the failed bank's deposits and will purchase $12.3 million of the failed bank's assets. To facilitate the transaction, the FDIC will advance about $49.8 million to the assuming bank and will retain assets of the failed

(more)

bank with a book value of about $62.0 million.

The Board of Directors approved the deposit assumption under its authority to do so whenever it determines that such a transaction will reduce the potential loss to the FDIC. The FDIC notes that its claim from the sale of the failed bank's assets will have priority over non-depositor creditors of the failed bank.

###

Appendix C. Model Derivations

This appendix provides the formal derivations for the structural model presented in Section IV. We derive the deposit franchise valuation, the asset valuation, the FDIC's scoring rule and the reduction of the multidimensional auction to a scalar first-price format, the optimal asset take-up share, the maximized private value, the comparative statics, and the sale-threshold condition for the probability of a P&A resolution.

Appendix C.A. Deposit Franchise Value

We derive the economic value $V_{ij}^{D}(\beta_i, \lambda_i)$ of failed bank i's deposit franchise to acquiring bank j.

Assumption C.A.1: *The core deposit base of failed bank i follows a deterministic exponential decay process from an initial stock $D_{i0} > 0$ at the auction date $t = 0$, governed by a constant hazard rate $\lambda_i > 0$:*

$$D_{it} = D_{i0} e^{-\lambda_i t}, \qquad t \geq 0. \tag{C.A.1}$$

Assumption C.A.2: *The market short-term interest rate r_t^m follows a stochastic process with constant unconditional mean $\mu_{r^m} := \mathbf{E}[r_t^m]$. Acquirer j earns a gross return on earning assets $r_{jt} > 0$ with mean $\mu_{r_j} := \mathbf{E}[r_{jt}]$ and incurs a constant marginal operating cost $c_j > 0$ per dollar of deposits serviced.*

Assumption C.A.3: *The deposit rate paid to customers of failed bank i adjusts linearly to the market rate:*

$$r_{it}^{d} = \alpha_i + \beta_i r_t^m, \tag{C.A.2}$$

where $\alpha_i \geq 0$ is a bank-specific intercept and $\beta_i \in [0, 1]$ is the deposit rate pass-through.

Assumption C.A.4: *The expected net interest margin is strictly positive:*

$$\mu_{r_j} - \alpha_i - \beta_i \mu_{r^m} - c_j > 0. \tag{C.A.3}$$

The instantaneous profit flow generated by the deposit base is

$$\Pi_t := \left(r_{jt} - r_{it}^{d} - c_j \right) D_{it}. \tag{C.A.4}$$

Substituting Assumptions C.A.1–C.A.3, the expected instantaneous profit is

$$\mathbf{E}[\Pi_t] = \left(\mu_{r_j} - \alpha_i - \beta_i \mu_{r^m} - c_j \right) D_{i0} e^{-\lambda_i t}. \tag{C.A.5}$$

The franchise value is the risk-neutral present value of this expected profit stream, discounted at

the acquirer's cost of capital $r_f > 0$:

$$V_{ij}^D(\beta_i, \lambda_i) := \int_0^\infty \mathbf{E}[\Pi_t] e^{-r_f t} dt \tag{C.A.6}$$

$$= \left(\mu_{r_j} - \alpha_i - \beta_i \mu_{r^m} - c_j\right) D_{i0} \int_0^\infty e^{-(r_f+\lambda_i)t} dt \tag{C.A.7}$$

$$= \frac{\left(\mu_{r_j} - \alpha_i - \beta_i \mu_{r^m} - c_j\right) D_{i0}}{r_f + \lambda_i}, \tag{C.A.8}$$

where the integral converges because $r_f + \lambda_i > 0$.

Appendix C.B. Asset Valuation

Assumption C.B.1: *The FDIC's expected recovery value per dollar of book assets for failed bank i is $r_i^{FDIC} \in (0, 1)$.*

Assumption C.B.2: *Bidder j's marginal valuation for the x-th fractional unit of bank i's asset book is linear and declining:*

$$v_{ij}(x) := \bar{v}_{ij} - \gamma_i x, \qquad \gamma_i > 0, \tag{C.B.1}$$

where $\bar{v}_{ij} \in (0, 1]$ is the marginal valuation for the first tranche and γ_i governs the rate of decline.

Let $s_{ij} \in [0, 1]$ denote the fraction of the asset book that bidder j proposes to acquire. The total asset value from acquiring share s_{ij} is obtained by integrating the marginal valuation:

$$V_{ij}^A(s_{ij}) := A_{i0} \int_0^{s_{ij}} v_{ij}(x) dx \tag{C.B.2}$$

$$= A_{i0} \int_0^{s_{ij}} \left(\bar{v}_{ij} - \gamma_i x\right) dx \tag{C.B.3}$$

$$= A_{i0} \left(\bar{v}_{ij} s_{ij} - \frac{1}{2} \gamma_i s_{ij}^2\right). \tag{C.B.4}$$

Appendix C.C. Scoring Rule and Reduction to a Scalar Auction

Each bidder j submits a sealed bid (p_{ij}, q_{ij}, s_{ij}) for failed bank i, consisting of a deposit premium rate $p_{ij} \geq 0$, an asset purchase price $q_{ij} \in [0, 1]$ per dollar of book, and an asset take-up share $s_{ij} \in [0, 1]$.

Assumption C.C.1: *The failed bank is insolvent in expectation: $D_{i0} > r_i^{FDIC} A_{i0}$.*

The net cost to the DIF under bid (p, q, s) is

$$C(p, q, s) = D_{i0} - (pD_{i0} + qsA_{i0}) - r_i^{\text{FDIC}}(1 - s)A_{i0}. \tag{C.C.1}$$

Rearranging:

$$C(p,q,s) = D_{i0} - r_i^{\text{FDIC}} A_{i0} - pD_{i0} - qsA_{i0} + r_i^{\text{FDIC}} sA_{i0} \tag{C.C.2}$$

$$= \underbrace{(D_{i0} - r_i^{\text{FDIC}} A_{i0})}_{\text{baseline shortfall } (>0)} - \underbrace{[pD_{i0} + (q - r_i^{\text{FDIC}}) sA_{i0}]}_{:= \mathcal{S}_i(p,q,s)}. \tag{C.C.3}$$

Under Assumption C.C.1 the baseline shortfall is positive and common across bids. The FDIC's least-cost mandate therefore reduces to maximizing the score

$$\mathcal{S}_i(p,q,s) := pD_{i0} + (q - r_i^{\text{FDIC}})\, sA_{i0}. \tag{C.C.4}$$

Assumption C.C.2: *There are $J \geq 2$ risk-neutral bidders. Each bidder j draws a private synergy parameter $\theta_{ij} \in (0,1]$ and an integration cost $\bar{K}_{ij} \in \mathbb{R}_+$, both unobserved by rivals and the FDIC. The deposit-franchise parameters (β_i, λ_i) and asset parameters $(\bar{v}_{ij}, \gamma_i)$ are common knowledge, though $\bar{v}_{ij}$ may vary across bidders.*

The payoff to bidder j upon winning with bid (p,q,s) is

$$U_{ij}(p,q,s) = \theta_{ij} V_{ij}^D(\beta_i, \lambda_i) + V_{ij}^A(s) - \bar{K}_{ij} - pD_{i0} - qsA_{i0}. \tag{C.C.5}$$

We now express the total payment in terms of the score. From (C.C.4):

$$pD_{i0} + qsA_{i0} = \mathcal{S}_i + r_i^{\text{FDIC}} sA_{i0}. \tag{C.C.6}$$

Substituting (C.C.6) and (C.B.2) into (C.C.5):

$$U_{ij} = \theta_{ij} V_{ij}^D + A_{i0}\left(\bar{v}_{ij} s - \frac{1}{2}\gamma_i s^2\right) - \bar{K}_{ij} - \mathcal{S}_i - r_i^{\text{FDIC}} sA_{i0} \tag{C.C.7}$$

$$= \underbrace{\theta_{ij} V_{ij}^D - \bar{K}_{ij} + A_{i0}\left[(\bar{v}_{ij} - r_i^{\text{FDIC}}) s - \frac{1}{2}\gamma_i s^2\right]}_{:= W_{ij}(s)} - \mathcal{S}_i. \tag{C.C.8}$$

The function $W_{ij}(s)$ is the bidder's *private value*: the surplus earned from winning at a score of zero. Equation (C.C.7) has the canonical structure of a first-price auction with private value W_{ij} and bid $\mathcal{S}_i$. The standard equilibrium analysis of Milgrom and Weber (1982) applies directly.

Appendix C.D. Optimal Asset Take-Up

Each bidder selects s to maximize $W_{ij}(s)$, which determines the "budget" for score bidding. Taking the first-order condition of $W_{ij}(s)$ with respect to s:

$$\frac{\partial W_{ij}}{\partial s} = A_{i0}\left(\bar{v}_{ij} - r_i^{\text{FDIC}} - \gamma_i s\right) = 0 \quad \Longrightarrow \quad s_{ij}^* = \frac{\bar{v}_{ij} - r_i^{\text{FDIC}}}{\gamma_i}. \tag{C.D.1}$$

The second-order condition confirms a global maximum:

$$\frac{\partial^2 W_{ij}}{\partial s^2} = -A_{i0}\gamma_i < 0. \tag{C.D.2}$$

Imposing the box constraint $s \in [0, 1]$:

$$s_{ij}^* = \min\left\{\frac{\bar{v}_{ij} - r_i^{\text{FDIC}}}{\gamma_i},\ 1\right\}, \tag{C.D.3}$$

with $s_{ij}^* \geq 0$ ensured whenever $\bar{v}_{ij} \geq r_i^{\text{FDIC}}$. The model-implied retained asset share under the winning bidder j^* is

$$\tau_i := 1 - s_{ij^*}^*. \tag{C.D.4}$$

Appendix C.E. Maximized Private Value

Substituting the interior solution $s_{ij}^* = (\bar{v}_{ij} - r_i^{\text{FDIC}})/\gamma_i$ into $W_{ij}(s)$:

$$W_{ij}(s_{ij}^*) = \theta_{ij} V_{ij}^D - \bar{K}_{ij} + A_{i0}\left[(\bar{v}_{ij} - r_i^{\text{FDIC}})\frac{\bar{v}_{ij} - r_i^{\text{FDIC}}}{\gamma_i} - \frac{1}{2}\gamma_i\left(\frac{\bar{v}_{ij} - r_i^{\text{FDIC}}}{\gamma_i}\right)^2\right] \tag{C.E.1}$$

$$= \theta_{ij} V_{ij}^D - \bar{K}_{ij} + A_{i0}\left[\frac{(\bar{v}_{ij} - r_i^{\text{FDIC}})^2}{\gamma_i} - \frac{(\bar{v}_{ij} - r_i^{\text{FDIC}})^2}{2\gamma_i}\right] \tag{C.E.2}$$

$$= \theta_{ij} V_{ij}^D(\beta_i, \lambda_i) - \bar{K}_{ij} + \frac{A_{i0}(\bar{v}_{ij} - r_i^{\text{FDIC}})^2}{2\gamma_i}. \tag{C.E.3}$$

The three terms are, respectively, the net deposit-side surplus, the integration cost, and the asset-side option value. The deposit and asset channels are additively separable: s_{ij}^* depends only on asset-side parameters $(\bar{v}_{ij}, \gamma_i, r_i^{\text{FDIC}})$, and V_{ij}^D depends only on deposit-side parameters $(\beta_i, \lambda_i, \mu_{r_j}, \alpha_i, c_j)$.

Appendix C.F. Symmetric Equilibrium Bidding

Assumption C.F.1: *The private values $W_{ij}(s_{ij}^*)$ are i.i.d. draws from a distribution G with support $[\underline{W}, \overline{W}]$, density $g > 0$, and CDF G. Each bidder submits a score bid $\mathcal{S}_i$ to maximize expected surplus.*

Define the normalized per-dollar private value $w_j := W_j/D_{i0}$, with CDF G and density g inherited from the distribution of W_j. The maximization problem for bidder j with value w is

$$\max_{\mathcal{S}} (w - \mathcal{S})\ \mathbf{P}\{\text{Win} \mid \mathcal{S}\}. \tag{C.F.1}$$

In a symmetric equilibrium with a strictly increasing bidding function $b(\cdot)$, the probability of

winning is the probability that all $J-1$ rivals have values below $b^{-1}(\mathcal{S})$:

$$\mathbf{P}\{\text{Win} \mid \mathcal{S}\} = G\left(b^{-1}(\mathcal{S})\right)^{J-1}. \tag{C.F.2}$$

Substituting $\mathcal{S} = b(w)$ and taking the first-order condition:

$$b'(w) = (J-1)\frac{g(w)}{G(w)}\left(w - b(w)\right). \tag{C.F.3}$$

This is a first-order linear ODE. With the boundary condition $b(\underline{w}) = \underline{w}$ (the lowest-type bidder earns zero surplus), the unique solution is

$$b(w) = w - \frac{\int_{\underline{w}}^{w} G(y)^{J-1} dy}{G(w)^{J-1}}. \tag{C.F.4}$$

Proof. Define $h(w) := b(w)G(w)^{J-1}$. Differentiating:

$$h'(w) = b'(w)G(w)^{J-1} + b(w)(J-1)G(w)^{J-2}g(w). \tag{C.F.5}$$

Substituting the first-order condition $b'(w)G(w)^{J-1} = (J-1)g(w)G(w)^{J-2}\left(w - b(w)\right)G(w)$:

$$\begin{aligned} h'(w) &= (J-1)g(w)G(w)^{J-2}\left(w - b(w)\right) + (J-1)g(w)G(w)^{J-2}b(w) && \text{(C.F.6)} \\ &= (J-1)g(w)G(w)^{J-2}w && \text{(C.F.7)} \\ &= w\frac{d}{dw}\left[G(w)^{J-1}\right]. && \text{(C.F.8)} \end{aligned}$$

Integrating from $\underline{w}$ to w with $h(\underline{w}) = \underline{w}G(\underline{w})^{J-1}$:

$$h(w) - h(\underline{w}) = \left[yG(y)^{J-1}\right]_{\underline{w}}^{w} - \int_{\underline{w}}^{w} G(y)^{J-1} dy, \tag{C.F.9}$$

where we integrate by parts. Simplifying:

$$b(w)G(w)^{J-1} = wG(w)^{J-1} - \int_{\underline{w}}^{w} G(y)^{J-1} dy, \tag{C.F.10}$$

which yields (C.F.4). □

The winning score observed in the data corresponds to the bid of the value leader $w^{(1)} := \max_j\{w_j\}$:

$$\mathcal{S}^* = b\left(w^{(1)}\right). \tag{C.F.11}$$

The bid-shading term $\int_{\underline{w}}^{w} G(y)^{J-1}dy / G(w)^{J-1}$ is strictly positive for $w > \underline{w}$, ensuring that the winning bidder earns a positive surplus.

Appendix C.G. Deposit-Side Comparative Statics

We derive the partial derivatives of the franchise value V_{ij}^D in (C.A.6) with respect to the deposit-side primitives β_i and λ_i.

Differentiating (C.A.6) with respect to β_i:

$$\frac{\partial V_{ij}^D}{\partial \beta_i}(\beta_i, \lambda_i) = \frac{-\mu_{r^m} D_{i0}}{r_f + \lambda_i} < 0. \tag{C.G.1}$$

The inequality follows from $\mu_{r^m} > 0$, $D_{i0} > 0$, and $r_f + \lambda_i > 0$.

Differentiating (C.A.6) with respect to λ_i:

$$\frac{\partial V_{ij}^D}{\partial \lambda_i}(\beta_i, \lambda_i) = -\frac{\left(\mu_{r_j} - \alpha_i - \beta_i \mu_{r^m} - c_j\right) D_{i0}}{(r_f + \lambda_i)^2} < 0. \tag{C.G.2}$$

The inequality follows from Assumption C.A.4.

Differentiating (C.G.1) with respect to λ_i, or equivalently (C.G.2) with respect to β_i:

$$\frac{\partial^2 V_{ij}^D}{\partial \beta_i \partial \lambda_i}(\beta_i, \lambda_i) = \frac{\mu_{r^m} D_{i0}}{(r_f + \lambda_i)^2} > 0. \tag{C.G.3}$$

Rate sensitivity and attrition are substitutes: the marginal cost of a higher β_i is decreasing in λ_i.

Appendix C.H. Asset-Side Comparative Statics

We derive the responses of the optimal take-up share s_{ij}^* in (C.D.3) and the maximized private value $W_{ij}(s_{ij}^*)$ in (C.E.3) to the asset-side primitives, restricting to interior solutions $s_{ij}^* < 1$.

Differentiating (C.D.3):

$$\frac{\partial s_{ij}^*}{\partial \bar{v}_{ij}} = \frac{1}{\gamma_i} > 0. \tag{C.H.1}$$

$$\frac{\partial s_{ij}^*}{\partial \gamma_i} = -\frac{\bar{v}_{ij} - r_i^{\text{FDIC}}}{\gamma_i^2} < 0. \tag{C.H.2}$$

$$\frac{\partial s_{ij}^*}{\partial r_i^{\text{FDIC}}} = -\frac{1}{\gamma_i} < 0. \tag{C.H.3}$$

Differentiating (C.E.3) with respect to $\bar{v}_{ij}$:

$$\frac{\partial W_{ij}}{\partial \bar{v}_{ij}}(s_{ij}^*) = \frac{A_{i0}(\bar{v}_{ij} - r_i^{\text{FDIC}})}{\gamma_i} = A_{i0} s_{ij}^* > 0, \tag{C.H.4}$$

where the second equality uses (C.D.3). An envelope-theorem argument confirms the same result: since s_{ij}^* maximizes $W_{ij}(s)$, the indirect effect through $\partial s^*/\partial \bar{v}_{ij}$ vanishes at the optimum, and only the direct effect $A_{i0} s_{ij}^*$ remains.

Differentiating (C.E.3) with respect to γ_i:

$$\frac{\partial W_{ij}}{\partial \gamma_i}(s_{ij}^*) = -\frac{A_{i0}(\bar{v}_{ij} - r_i^{\text{FDIC}})^2}{2\gamma_i^2} < 0. \tag{C.H.5}$$

Differentiating (C.E.3) with respect to r_i^{FDIC}:

$$\frac{\partial W_{ij}}{\partial r_i^{\text{FDIC}}}(s_{ij}^*) = -\frac{A_{i0}(\bar{v}_{ij} - r_i^{\text{FDIC}})}{\gamma_i} = -A_{i0}s_{ij}^* < 0. \tag{C.H.6}$$

Appendix C.I. Premium–Retention Relationship

We establish the positive association between the retained share τ_i and the deposit premium p through two channels.

From (C.C.4), the deposit premium consistent with a given score $\mathcal{S}_i$ and asset terms (q, s) is

$$p = \frac{\mathcal{S}_i - (q - r_i^{\text{FDIC}})sA_{i0}}{D_{i0}}. \tag{C.I.1}$$

Fix $\mathcal{S}_i$ and q. Then $\partial p/\partial s = -(q - r_i^{\text{FDIC}})A_{i0}/D_{i0}$. At the optimal take-up, $q = v_{ij}(s_{ij}^*) = r_i^{\text{FDIC}}$ at the margin, and for infra-marginal units $q > r_i^{\text{FDIC}}$, so $(q - r_i^{\text{FDIC}}) \geq 0$. A lower s (higher τ_i) therefore weakly increases p for a given score.

From (C.E.3), the maximized private value is increasing in the comparative-advantage wedge $\bar{v}_{ij} - r_i^{\text{FDIC}}$. When marginal assets satisfy $v_{ij}(x) < r_i^{\text{FDIC}}$, retaining them (increasing τ_i) removes a drag on W_{ij}, raising the equilibrium score bid through the monotonicity of $b(\cdot)$ in (C.F.4). The higher score translates into a higher premium via (C.I.1). Both channels sign the τ_i–p relationship as positive.

Appendix C.J. Probability of P&A

Let C_i^{payoff} denote the DIF cost under a deposit payoff. A P&A sale occurs if and only if

$$\max_j \ \mathcal{S}_i(p_{ij}, q_{ij}, s_{ij}) \ \geq \ \underbrace{(D_{i0} - r_i^{\text{FDIC}}A_{i0}) - C_i^{\text{payoff}}}_{:=\underline{\mathcal{S}}_i}. \tag{C.J.1}$$

In the symmetric equilibrium, the score bid $b(w)$ is strictly increasing. The event {sale} is therefore equivalent to $\{w^{(1)} \geq b^{-1}(\underline{\mathcal{S}}_i)\}$. Let $\underline{w}_i := b^{-1}(\underline{\mathcal{S}}_i)$ denote the threshold private value. The probability of a P&A is

$$\mathbf{P}\{\text{P\&A}\} = 1 - G(\underline{w}_i)^J. \tag{C.J.2}$$

We verify three comparative statics:

(i) *Franchise value.* An increase in V_{ij}^D shifts the distribution of W_{ij} to the right (by (C.E.3)), increasing $G(\underline{w}_i)^J$'s complement.

(ii) *Asset-side comparative advantage.* An increase in $\bar{v}_{ij} - r_i^{\text{FDIC}}$ raises the asset-side option value in (C.E.3), shifting W_{ij} to the right.

(iii) *Number of bidders.* Differentiating (C.J.2) with respect to J:

$$\frac{\partial}{\partial J}\left[1 - G(\underline{w}_i)^J\right] = -G(\underline{w}_i)^J \ln G(\underline{w}_i) > 0, \tag{C.J.3}$$

since $G(\underline{w}_i) \in (0,1)$ implies $\ln G(\underline{w}_i) < 0$.

All three channels increase the probability that the maximum score bid exceeds the least-cost threshold $\underline{\mathcal{S}}_i$, completing the derivation of the model's testable predictions.

Appendix D. Robustness

Appendix D.A. IV First-Stage Regression

	Toxic Ratio
LLP Ratio	0.89711***
	(0.23143)
Log Asset Size	0.00691
	(0.00599)
Year FE	X
N	1785
R^2	0.25
First-stage F	15.0

Table X: OLS regression of the realized toxic ratio on the pre-failure loan loss provision ratio and controls. The instrument is the loan loss provision ratio (LLP / assets) in the quarter before failure. Year fixed effects included. F-statistic for instrument relevance reported. Robust standard errors in parentheses. Significance levels: $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Appendix D.B. Heckman Selection-Corrected Pricing

This appendix provides the complete econometric derivation, first-stage probit estimates, bias decomposition, and robustness analysis for the Heckman two-step correction summarized in Section V.

Table XI reports the first-stage probit estimates. The dependent variable is the P&A indicator. Column (1) presents the primary specification, which includes deposit runoff, the LLP Ratio, log asset size, and the quarterly concurrent-failure exclusion variable with year fixed effects. Column (2) replaces the congestion measure with log days elapsed since the nearest macroeconomic crisis trough. Column (3) includes both exclusion variables simultaneously to form an over-identified specification. Column (4) restricts to the deposit-beta matched subsample and adds deposit beta as an additional selection covariate.

The probit estimates are consistent with the linear probability model of Section V. Asset size enters with a large positive coefficient in all specifications: larger institutions are substantially more likely to attract a P&A acquirer. Deposit runoff carries a negative point estimate that is imprecisely estimated once year effects absorb the macroeconomic flight component. The LLP Ratio enters positively across all columns, consistent with the asset-cleansing narrative: banks for which a larger loss-sharing pool has been negotiated are precisely those for which sufficient acquirer interest exists to complete a full P&A.

The coefficient on log concurrent failures in Column (1) is positive and marginally significant. This sign is the opposite of the ex ante capacity-constraint prediction, but it is consistent with the

	$\mathbf{1}\{\text{Resolution}_i = \text{P\&A}\}$			
	(1)	(2)	(3)	(4)
Deposit Runoff	-0.46428	-0.44134	-0.43073	0.26531
	(0.40928)	(0.41929)	(0.42101)	(1.75178)
LLP Ratio	1.02656	0.98665	1.01374	1.90658
	(1.17418)	(1.17059)	(1.17450)	(7.14635)
Log Asset Size	0.03607	0.03733	0.03573	-0.00374
	(0.03437)	(0.03427)	(0.03440)	(0.10845)
concurrent_failures	0.29254*		0.29006*	0.33925
	(0.15959)		(0.15980)	(0.41185)
cr_gap_log		-0.02524	-0.02030	
		(0.05922)	(0.05927)	
Year FE	X	X	X	X
N	1784	1784	1784	372
Pseudo R^2	0.098	0.096	0.098	0.110
Log-likelihood	-723.0	-724.6	-722.9	-50.1

Table XI: Probit regressions of an indicator for P&A resolution on deposit and asset controls plus exclusion restrictions. Column (1) uses log concurrent failures in the resolution quarter; column (2) uses log days from the nearest crisis trough; column (3) includes both exclusion variables; column (4) restricts to the deposit-beta matched subsample and adds the deposit expense beta. Probit marginal index coefficients. Robust standard errors in parentheses. Significance $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

FDIC's pro-cyclical resolution infrastructure: when failures cluster in a single quarter, fixed costs of assembling bidder lists and coordinating with chartering authorities are amortized across many institutions, and the agency has stronger incentives to clear the resolution queue quickly through P&A rather than payoff. This pattern is consistent with the historical narrative of the S&L crisis resolution documented in Curry and Shibut (2000). For the identification argument, the sign of the exclusion coefficient is immaterial; what matters is that the variable enters the selection equation with a non-zero coefficient and is plausibly absent from the conditional pricing equation.

In Column (4), deposit beta enters with a negative coefficient, reinforcing the interpretation that banks with greater deposit rate pass-through are less marketable in a resolution auction.

Appendix D.C. Bias Decomposition

Table XII consolidates the selection-bias diagnostics across the four Heckman specifications.

Three substantive points follow from the decomposition. First, $\hat{\sigma}_\varepsilon$ is stable across specifications at approximately 0.015, matching the sample standard deviation of the log deposit premium reported in Section V. Second, the implied correlation $\hat{\rho}$ is moderate in the primary specification at 0.42 but is sensitive to the choice of exclusion restriction and even reverses sign when the alternative exclusion is used. The corresponding bias in the pricing coefficients is bounded by

	(1) Concurrent	(2) Crisis Gap	(3) Over-Id	(4) β
$\hat{\rho}$	0.419	-0.207	0.286	1.794
$\hat{\sigma}_\varepsilon$	0.015	0.015	0.015	0.017
$\hat{\beta}_\lambda$	0.006	-0.003	0.004	0.030
LR test (over-id)			0.12	
N (first stage)	1784	1784	1784	372
N (second stage)	1488	1488	1488	359

Table XII: The column entries report the implied correlation $\hat{\rho}$ between the pricing and selection errors, the second-stage residual standard deviation $\hat{\sigma}_\varepsilon$, the fitted inverse Mills ratio coefficient $\hat{\beta}_\lambda = \hat{\rho}\hat{\sigma}_\varepsilon$, and the first-stage log-likelihood ratio test statistic for exclusion-restriction validity. $\hat{\rho}$ is estimated as $\hat{\beta}_\lambda/\hat{\sigma}_\varepsilon$ where $\hat{\sigma}_\varepsilon^2 = \hat{u}'\hat{u}/n + \hat{\beta}_\lambda^2\bar{\delta}$ with $\bar{\delta} = n^{-1}\sum_i \hat{\lambda}_i(\hat{\lambda}_i + \hat{x}_i\hat{\gamma})$.

$|\hat{\rho}|\hat{\sigma}_\varepsilon\bar{\lambda} \approx 0.42 \cdot 0.015 \cdot 1.3 \approx 0.008$, an order of magnitude comparable to the standard errors on the pricing coefficients themselves. Third, the likelihood-ratio statistic for the joint significance of the two exclusion variables in the over-identified probit is 0.12 on two degrees of freedom, well below any conventional critical value. The two exclusion variables are effectively one-dimensional in their first-stage contribution, capturing overlapping aspects of resolution-market conditions. The mild over-identification failure does not invalidate the identification strategy because the pricing coefficients are robust across the single-exclusion specifications.

Column (3) of Table XI includes both log concurrent failures and log days from crisis trough simultaneously. The two exclusion variables capture different sources of variation: the former reflects within-year seasonal congestion in resolution volume; the latter reflects the position of any given failure relative to the aggregate macroeconomic cycle. Despite this difference, the pseudo-R^2 of the over-identified probit rises only modestly relative to the single-exclusion specifications, and the pricing coefficients in Column (4) of Table VIII are effectively identical to those in the primary specification. The likelihood-ratio over-identification test statistic of 0.12 (2 df) implies that neither exclusion variable dominates the other in first-stage identification. The practical implication is that identification of the Heckman correction rests on a single effective instrument. That instrument is marginally significant (first-stage $t \approx 1.83$), which is part of the reason for the wide bootstrap confidence interval on $\hat{\beta}_\lambda$ and motivates emphasizing coefficient stability across specifications rather than point-estimate precision.-

Column (5) of Table VIII reports the Heckman second stage on the deposit-beta matched subsample. The implied $\hat{\rho}$ in this column lies outside the unit interval, a telltale symptom of finite-sample Heckman pathology documented by Puhani (2000). The pathology arises because the deposit-beta subsample is small, selected on survival (a bank must operate long enough to estimate a rate-sensitivity beta), and exhibits limited variation in the inverse Mills ratio, reducing the effective identification of β_λ. This column should be interpreted as a diagnostic robustness check rather than a refined estimate of the selection correction.

The Heckman correction and the 2SLS strategy of Section V address different identification threats: the former targets selection into the P&A sample, the latter targets the simultaneity of

the premium and the retained asset share. A fully identified specification would combine both corrections in a selection-corrected instrumental-variable design. I have estimated this combined specification and find pricing coefficients that are economically identical to those reported in Tables V and VIII, with larger standard errors reflecting the compounding of generated-regressor uncertainty across the two stages. The separate corrections are reported throughout to preserve the diagnostic interpretation of each. The dominant identification threat in the pricing channel remains the simultaneity of the toxic ratio and the premium, not the selection into the P&A sample.

Several robustness margins merit brief discussion. Redefining log concurrent failures at the monthly or semi-annual frequency, or replacing the log transformation with the level count, does not materially affect the first-stage significance or the pricing coefficients. Restricting the sample to the post-FIRREA period ($N \approx 1{,}000$) weakens first-stage identification further, as the most congested quarters of the sample fall in the thrift crisis; the pricing coefficients remain within the bootstrap confidence intervals reported in Table VIII.

The joint-normality assumption of the parametric Heckman model is routinely questioned on the grounds that non-normality produces an inconsistent control function. I have re-estimated the pricing equation using a semi-parametric single-index selection correction following Newey (2009) on the primary exclusion restriction. The resulting pricing coefficients are within one bootstrap standard error of the parametric two-step estimates, indicating that the conclusions of Section V are not driven by the normality assumption.

Restricting the selected sample to pure P&A transactions (excluding Insured Deposit Transfers) or expanding it to include all transactions involving any form of franchise transfer does not alter the finding that the selection correction leaves the pricing coefficients materially unchanged.

REFERENCES

102nd Congress of the United States, 1991, Federal deposit insurance corporation improvement act of 1991.

Barth, James R., and Apanard Prabha, 2013, Resolving too-big-to-fail banks in the united states, *Mercatus Center Working Paper Series* 13.

Bennett, Rosalind L., and Haluk Unal, 2015, Understanding the components of bank failure resolution costs, *Financial Markets, Institutions & Instruments* 24, 349–389.

Benston, George J., and George G. Kaufman, 1997, Fdicia after five years, *Journal of Economic Perspectives* 11, 139–158.

Bliss, Robert R., and George G. Kaufman, 2006, A comparison of u.s. corporate and bank insolvency resolution, *Federal Reserve Bank of Chicago: Economic Perspectives* 30.

Bovenzi, John F., and Maureen E. Muldoon, 1990, Failure-resolution methods and policy considerations, *FDIC Banking Review* 3, 1–11.

Bovenzi, John F., and Arthur J. Murton, 1988, Resolution costs of bank failures, *FDIC Banking Review* 1, 1–13.

Brown, Richard A., and Seth Epstein, 1992, Resolution costs of bank failures: An update of the fdic historic al loss model, *FDIC Banking Review* 5, 1–16.

Corporation, Federal Deposit Insurance, 2003, Failed bank cost analysis: 1986-2002, *FDIC Publication* .

Cowan, Arnold R., and Valentina Salotti, 2015, The resolution of failed banks during the crisis: Acquirer performance and fdic guarantees, 2008-2013, *Journal of Banking & Finance* 54, 222–238.

Cowan, Arnold R., Valentina Salotti, and Natalya A. Schenck, 2022, The long-term impact of bank mergers on stock performance and default risk: The aftermath of the 2008 financial crisis, *Finance Research Letters* 48, 102925.

Curry, Timothy, and Lynn Shibut, 2000, The cost of the savings and loan crisis: Truth and consequences, *FDIC Banking Review* .

Davison, Lee K., 1997, *Continental Illinois and Continental Illinois and "Too Big to F oo Big to Fail"*, chapter 7, 236–258 (Federal Deposit Insurance Corporation).

Diamond, Douglas W., and Philip H. Dybvig, 1983, Bank runs, deposit insurance, and liquidity, *Journal of Political Economy* 91, 401–419.

Drechsler, Itamar, Alexi Savov, and Philipp Schnabl, 2017, The deposits channel of monetary policy*, *The Quarterly Journal of Economics* 132, 1819–1876.

Drechsler, Itmar, Alexi Savov, and Philipp Schnabl, 2021, Banking on deposits: Maturity transformation without interest rate risk, *The Journal of Finance* 76, 1091–1143.

Federal Deposit Insurance Corporation, and Resolution Trust Corporation, 1998, *Managing the Crisis: The FDIC and RTC Experience, 1980–1994* (Federal Deposit Insurance Corporation, Washington, DC), 2 vols.

Granja, Joao, Gregor Matvos, and Amit Seru, 2017, Selling failed banks, *The Journal of Finance* 72, 1723–1784.

Heckman, James J., Hidehiko Ichimura, and Petra E. Todd, 1997, Matching as an econometric evaluation estimator: Evidence from evaluating a job training programme, *The Review of Economic Studies* 64, 605–654.

Heitz, Amanda, 2022, The long-run impact of government asset guarantees, *FDIC Working Paper Series* .

Hirschhorn, Eric, and David Zervos, 1990, Policies to change the priority of claimants: The case of depositor preference laws, *Journal of Financial Services Research* 4, 111–125.

James, Christopher, 1991, The losses realized in bank failures, *The Journal of Finance* 46, 1223–1242.

Marino, James A, and Rosalind L. Bennett, 1999, The consequences of national depositor preference, *FDIC Banking Review* 12, 19–38.

McCoy, Patricia A., 2007, *The Moral Hazard Implications of Deposit Insurance: Theory and Evidence*, chapter 16, 417–442 (International Monetary Fund).

McDill, Kathleen, 2004, Resolution costs and the business cycle, *FDIC Working Paper Series* .

Milgrom, Paul R., and Robert J. Weber, 1982, A theory of auctions and competitive bidding, *Econometrica* 50, 1089–1122.

Molyneux, Philip, and Tim Mi Zhou, 2022, Banking market reaction to auctions of failed banks, *International Journal of Finance & Economics* 27, 518–534.

Morgan, Donald P., Maxim L. Pinkovskiy, and Bryan Yang, 2016, Banking deserts, branch closings, and soft information, Liberty Street Economics 20160307, Federal Reserve Bank of New York.

Newey, Whitney K., 2009, Two-step series estimation of sample selection models, *The Econometrics Journal* 12, S217–S229.

Osterberg, William P., 1996, The impact of depositor preference laws, *Economic Review* None, 2–11.

Puhani, Patrick, 2000, The heckman correction for sample selection and its critique, *Journal of Economic Surveys* 14, 53–68.

Spokeviciute, Laima, Hossein Jahanshahloo, Kevin Keasey, and Francesco Vallascas, 2025, Three decades of failed bank acquisitions, *Journal of Banking and Finance* 170.

Staiger, Douglas, and James H. Stock, 1997, Instrumental variables regression with weak instruments, *Econometrica* 65, 557–586.