

7534498-the-financing-of-debt-relief-intermediary-liquidity-and-the-value-of-borrower-in

SUMMARY: This article examines how intermediaries allocate scarce cash to deliver debt relief in the mortgage servicing market. It focuses on Ginnie Mae buyout decisions from 2019 to 2022, analyzing the behavior of nonbank issuers and depositories. The study reveals that nonbank issuers cut exercise almost fivefold while depositories raised it, indicating that information and liquidity interact to determine which borrowers receive relief. The authors propose a model to understand this interaction, which ranks resolution opportunities by value per dollar of scarce cash. They find that nonbanks' response to the ranking remained close to pre-shock levels, even as funding returned. The study also shows that nonbanks postponed some exercise in 2020 but quickly resumed it, and that scarce cash shifted the ranking price, tilting exercise against loans with larger balances. These findings have implications for understanding how intermediaries allocate resources and the value of borrower information in the mortgage servicing market.

METHOD: The study uses 1,922,668 Ginnie Mae buyout decisions by 120 issuers from January 2019 to September 2022. The researchers analyze nonbank and depository behavior during a liquidity shock in 2020 and its reversal in 2021. The model ranks resolution opportunities by value per dollar of scarce cash, with a liquidity shock moving the ranking down and a value shock reordering it. The authors estimate nonbanks' ranking of loans from their own decisions in 2016-2018, within issuer-months to absorb budgets.

KEY FINDINGS:

- Nonbank issuers cut exercise almost fivefold while depositories raised it, indicating that information and liquidity interact to determine which borrowers receive relief.
- The response of nonbank exercise to the ranking relative to the base rate stayed close to its pre-shock level, even as funding returned.
- Most of the exercise nonbanks postponed in 2020 was made up within two years, and later exercise weighted the ranking about seven-tenths as heavily as exercise had before the shock.
- Scarcely available cash changed the ranking price, tilting exercise against loans with larger balances.

IMPLICATION FOR TRADING: The findings have implications for understanding how intermediaries allocate resources and the value of borrower information in the mortgage servicing market. Practitioners can use this information to better understand how intermediaries prioritize loan resolution and the value of borrower information in their decision-making. This can help in predicting future trends and optimizing resource allocation in the mortgage servicing market.

The Financing of Debt Relief:

Intermediary Liquidity and the Value of Borrower Information

Arka P. Bandyopadhyay*

September 2026

Latest version: arkabandyopadhyay.com

Abstract

Intermediaries deliver debt relief by spending cash before its value arrives. In 1.9 million Ginnie Mae buyout decisions from 2019 to 2022, nonbank issuers, whose funding tightened in 2020, cut exercise almost fivefold while depositories raised it, tilted against large balances while cash was scarce, and worked down postponed options along their pre-shock ranking as funding returned; when rates rose in 2022, depositories began to screen. A model of rationing by value per dollar organizes them: information and liquidity are complements when funding is scarce and substitutes when it is abundant, so AUC can misrank models for constrained servicers.

JEL classification: G21, G23, G28, G51, D82, D83.

Keywords: debt relief; mortgage servicing; intermediary liquidity; value of information; machine learning; Ginnie Mae.

*Department of Economics, Columbia University; Department of Economics, New York University; and Department of Economics, Yeshiva University. Email: ab3985@columbia.edu. Phone: (347) 260-0869. I am grateful to Patrick Smith and Dongshin Kim for sharing the Ginnie Mae loan-level data they assembled and for their guidance and mentorship, from which I have benefited significantly. The paper builds on companion work on the Ginnie Mae buyout market and on servicers' relief decisions (Bandyopadhyay, 2026a,b,c,d) and on Bandyopadhyay, Kim, and Smith (2026). Code reproducing every figure, every numerical table, and every numerical result accompanies the paper. Declarations of interest and of the use of AI-assisted technologies appear before the references. All errors are my own.

1 Introduction

Congress legislated mortgage relief in 2020, but intermediaries had to pay for it. A Ginnie Mae issuer must keep advancing scheduled principal and interest to bondholders on every loan whose borrower stops paying. Its one contractual exit, repurchasing the loan out of the pool at par once the borrower has missed three payments, requires the full unpaid balance in cash. Between the year before the pandemic and March through June 2020, the share of newly vested buyout options that nonbank issuers exercised fell from 40.0% to 8.3%, while depositories' rose from 80.5% to 85.3%. The option had not lost its value. Mortgage rates fell to record lows in 2020, which raised the premium at which a repurchased loan could be redelivered once it reperformed. Issuers are the parties best placed to know which delinquent borrowers will recover (Bandyopadhyay, Kim, and Smith, 2026). Many of them could not fund acting on what they knew.

The episode illustrates a general feature of how debt relief is delivered. Advancing payments, buying a loan out of a pool, capitalizing arrears into a modification, and staffing a loss-mitigation pipeline all require the intermediary to spend cash before the value of resolution arrives, so the intermediary's balance sheet is part of the relief technology. So is its information about borrowers. This paper asks how the two interact: which borrowers receive relief when liquidity tightens, what borrower information is worth to an intermediary that cannot fund every action its information recommends, and how such an intermediary should evaluate the predictive models it relies on.

I answer these questions with a model of an intermediary that allocates scarce cash across resolution opportunities, and with 1,922,668 buyout decisions by 120 Ginnie Mae issuers from January 2019 to September 2022. The period contains a liquidity shock concentrated on nonbank issuers, its reversal in 2021, and two shocks to the value of repurchase, and the model says which statistics each kind of shock should move. In the model, actions are ranked by value per dollar of scarce cash; a liquidity shock moves how far down its ranking a constrained intermediary goes and leaves the ranking alone, while a value shock moves

unconstrained intermediaries and can reorder the ranking for everyone. I estimate nonbanks' ranking of loans from their own decisions in 2016–2018, within issuer-months so that budgets are absorbed, and follow their use of that ranking through the cycle.

Four findings emerge. First, the funded share collapsed at nonbanks and rose at depositories, and the response of nonbank exercise to the ranking relative to the base rate, which random rationing would have lowered slightly and rationing along a fixed ranking would have raised sharply, stayed close to its pre-shock level; in issuer-by-issuer estimates, the largest nonbanks and the small issuers as a group kept responding to the ranking, and one large issuer whose funded share fell to 2.4% did not. As funding returned in the second half of 2020, the response rose to 0.67, close to the fixed-ranking value of 0.71 and above the random-rationing benchmark of 0.47, and by 2021 it had returned to where the pre-shock weight and the recovered funded share put it. Second, most of the exercise nonbanks postponed in 2020 was made up within two years, and later exercise weighted the ranking about seven-tenths as heavily as exercise had before the shock: the issuers had not stopped discriminating; they had stopped paying.

Third, scarce cash changed the price of cash in the ranking. Proposition 1 ranks by value per dollar, so a binding budget should tilt exercise against loans with larger balances holding value fixed. Nonbanks' weight on log balance, slightly positive before the shock, fell by 0.32 log-odds per standard deviation (s.e. 0.06) in the freeze and returned to its pre-shock level as funding came back, while depositories' did not move; none of the alternatives to rationing that the paper considers predicts this. Within issuer types, a nonbank's own stock of advances on delinquent loans lowers its exercise within issuer and month, with a smaller and insignificant response at depositories, but the sector-wide collapse of 2020 dwarfs the cross-issuer variation in this measure, so the funding channel is identified mainly from the contrast between types, the tilt, and the dynamics. Fourth, value shocks look different from liquidity shocks. When mortgage rates rose in 2022, depositories, whose funding the 2020 evidence suggests did not bind, cut exercise from 81.0% to 51.9% and began to use the ranking to decide which options

to decline, consistent with the switch from ranking value to screening value that the model predicts near full funding. In May and June 2020, depositories exercised 87.6% of vested options on loans in forbearance and nonbanks 3.9%, a gap that remains within issuer-months and within deciles of the ranking; a multi-period budget explains why constrained issuers left aside the repurchases that tie up cash longest.

The model delivers these predictions from one structure. An intermediary funds every opportunity whose value per dollar of cash exceeds a shadow price, so profitable relief goes unfunded whenever that price is positive, fewer borrowers receive relief when liquidity tightens, and the relief that is delivered performs better because the marginal recipients are the first to be dropped: a contraction in access looks like an improvement in measured performance. Under a binding budget, the cash committed to relief responds one-for-one to liquidity and is locally invariant to shocks to the value of relief; an unconstrained intermediary does the reverse. This double dissociation is what the type contrast in 2020 and the value shocks of 2020 and 2022 test.

The central theoretical result concerns information. The allocation value of a borrower signal is hump-shaped in liquidity: zero when nothing can be funded, rising and then falling as liquidity grows, and at full funding equal to a pure screening value, the value of declining actions that the signal reveals to be unprofitable. The peak occurs where the fundable share equals the share of borrowers whose signal is good news, and its height is one-half of the payoff scale times the mean absolute deviation of the posterior. Information and liquidity are therefore complements when liquidity is scarce and substitutes when it is abundant, and the same structure governs how strongly decisions respond to a ranking variable: the level response is hump-shaped in liquidity, while the response relative to the base rate rises as liquidity contracts. For model evaluation, a score's ranking value at a given liquidity is proportional to the vertical distance between its cumulative accuracy profile and the diagonal at the fundable share; one score dominates another at every liquidity level if and only if its posteriors dominate in the convex order; and the area under the ROC curve is an affine

function of allocation value averaged over liquidity levels drawn uniformly from $[0, 1]$. AUC therefore ranks models by their average over balance sheets the servicer does not have, and the ranking can reverse at the liquidity a particular servicer actually has. Two further results, on when information relaxes the funding constraint and on borrower welfare and policy, close the paper: information expands funding only when lenders can verify the technology that produces it, private selection can lower borrower surplus even as it raises private value, and ex post compensation cannot raise the cash a constrained intermediary commits to relief.

Contribution. The paper contributes to the literature on frictions in mortgage relief. Securitization, servicer incentives, and organizational capacity limit modification (Piskorski, Seru, and Vig, 2010; Adelino, Gerardi, and Willen, 2013; Agarwal et al., 2017; Maturana, 2017; Kruger, 2018). Servicer liquidity is a distinct friction: Kim et al. (2018) describe how advancing obligations expose nonbank servicers to liquidity crises, Aiello (2022) shows that financially constrained servicers reduce investor value through the modifications they choose, and Kim et al. (2024) show that servicer frictions shaped the delivery of CARES Act forbearance (see also Cherry et al., 2021); the growth of nonbank intermediaries without deposit funding makes these constraints first order (Buchak et al., 2018; Jiang et al., 2020). These papers establish that constrained servicers deliver less relief and different relief. This paper asks what a constrained servicer does with what it knows. It models funding and information jointly, shows that their interaction determines who receives relief and what borrower information is worth, and tests the mechanism on a decision where the cash requirement is explicit and the ranking can be recovered from the intermediary’s own choices. It also differs from the closest companion work. Bandyopadhyay (2026d) documents the 2020 collapse of nonbank exercise, its difference-in-differences and within-pool estimates, and the reversal of size gradients, on data through September 2020. The present paper takes that collapse as given and studies its structure: whether constrained issuers rationed along a fixed ranking as budgets contracted and recovered through 2022, what happened to the options

they postponed, how the same issuers responded to value shocks, and what all this implies for the value of information and the evaluation of predictive models. None of those results appears in the companion paper, and the theory is new.

The paper also relates to theories of intermediary capital and internal capital allocation (Holmström and Tirole, 1997, 1998; Stein, 1997, 2002; Froot, Scharfstein, and Stein, 1993), to work on hard and soft information (Liberti and Petersen, 2019), to the effect of financing constraints on the timing of real options (Boyle and Guthrie, 2003), and to the pledgeability of information (cf. Myers and Majluf, 1984; DeMarzo and Duffie, 1999). The value-of-information results use the decision-theoretic ordering of Blackwell (1953) and convex-order characterizations (Shaked and Shanthikumar, 2007); the contribution is to show that the budget determines which features of an information structure matter. Finally, the paper speaks to the literature on machine learning in credit markets, where better prediction redistributes access (Fuster et al., 2022), new data improve default prediction (Berg et al., 2020), technology changes mortgage lending (Fuster et al., 2019), and algorithmic decisions face selective labels (Kleinberg et al., 2018). In machine learning, AUC has been criticized as an incoherent measure of classifier performance (Hand, 2009) and corresponds to expected loss under a uniform distribution of selection rates (Hernández-Orallo, Flach, and Ferri, 2012); I provide the economic counterpart for financial intermediaries, in which the relevant distribution is the intermediary’s liquidity.

Section 2 describes the institutional setting. Section 3 develops the rationing model. Section 4 contains the information results. Section 5 studies borrowers, stakeholders, and policy. Section 6 presents the evidence and sets out the remaining tests. Section 7 concludes. Proofs are in Appendix A; the Internet Appendix contains auxiliary results, a dynamic extension, a data-construction protocol, the empirical details and robustness, and the numerical verification.

2 Institutional Setting

2.1 Cash-intensive relief

A Ginnie Mae issuer pools government-insured mortgages into securities and guarantees timely payment to investors. When a borrower misses a payment, the issuer advances scheduled principal and interest from its own funds for as long as the loan remains delinquent and in the pool. Once the borrower has missed three payments, the issuer may repurchase the loan at par. Repurchase ends the advance obligation on that loan. It also requires the full unpaid balance in cash and places a nonearning asset on the issuer’s balance sheet until the loan is cured, modified, or liquidated. A repurchased loan that reperforms can be redelivered into a new pool, which is valuable when market rates are below the loan’s rate. The option to repurchase belongs to the issuer of record and cannot be sold separately from the servicing (Bandyopadhyay, 2026d). Depositories fund these outlays from deposits. Nonbanks, which by 2020 serviced most Ginnie Mae collateral, fund them from capped warehouse and servicing-advance facilities (Kim et al., 2018; Jiang et al., 2020; Bandyopadhyay, 2026d).

A repurchase is not itself debt relief. It changes who holds the loan and the holder’s incentives without changing the borrower’s obligations. Relief comes from forbearance, deferral, or modification, each with its own cash requirement, compensation, and consequences for the borrower. The CARES Act of March 27, 2020, entitled borrowers with federally backed mortgages to forbearance on request, which multiplied delinquencies, and therefore advance obligations, across every issuer at once. Private-label servicers faced no statutory mandate but made discretionary relief decisions under their own contracts and funding (Bandyopadhyay, 2026c). The framework below treats these actions as different uses of the same scarce resource.

2.2 Two interventions of different kinds

Two 2020 interventions illustrate the distinction the model draws between liquidity and value. On April 10, 2020, Ginnie Mae made its Pass-Through Assistance Program available to issuers facing pandemic-related shortfalls (Ginnie Mae, 2020a). The program advanced funds to cover principal and interest owed to investors at a fixed interest rate, required full repayment, and was intended as an extraordinary measure for issuers whose other resources were exhausted. It was liquidity at a steep marginal cost, not a transfer. On June 29, 2020, Ginnie Mae restricted the redelivery of reperforming loans (Ginnie Mae, 2020b). A loan that entered forbearance on or after March 1, 2020, and was repurchased on or after July 1, 2020, became ineligible for existing pool types. It could be securitized again only in a new pool type, after six consecutive timely payments and at least 210 days after its last delinquency. The restriction lowered the value of repurchasing a dated subset of loans without supplying or absorbing any cash. In the model's terms, the first intervention relaxes the budget at a cost and the second shifts the value of action. A market-wide shift in value followed in 2022, when mortgage rates more than doubled and the redelivery value of most delinquent loans fell.

2.3 Four values that must be kept distinct

For each action I distinguish four quantities, all measured relative to the same feasible continuation. *Private incremental value* is the discounted payoff to the intermediary, including contractual transfers and financing at normal terms. *Cash required now* is the net outlay needed to implement the action before its proceeds arrive. *Borrower benefit* is the incremental effect on the borrower's payments, housing retention, and resources. *Total surplus* adds the gains of all parties net of real resource and fiscal costs, counting transfers once. Principal paid by an issuer to bondholders at repurchase is a private cash outflow and a transfer between stakeholders, not a social loss. Conversely, a loan that reperforms after an intervention can be privately valuable without any borrower gain if the borrower would have cured anyway.

The model keeps these objects separate because the policy conclusions turn on the differences among them.

3 Rationing Relief

3.1 Primitives

A risk-neutral intermediary holds a unit mass of resolution opportunities indexed by i , with distribution F . For each opportunity, it chooses between a normalized continuation and an intervention, $x_i \in [0, 1]$; interior values represent randomization across otherwise identical opportunities. The intervention has private incremental value g_i relative to continuation and requires net cash $\ell_i > 0$ at the decision date. The value g_i already includes the present value of all future receipts and costs. The cash requirement appears separately because cash cannot be moved freely across time. Liquidity $K \geq 0$ is usable cash and committed borrowing capacity net of unavoidable obligations. The intermediary solves

$$V(K) = \sup_{0 \leq x_i \leq 1} \int g_i x_i dF(i) \quad \text{subject to} \quad \int \ell_i x_i dF(i) \leq K. \quad (1)$$

Assumption 1. *The population is nonatomic; g is bounded and measurable; and $0 < \underline{\ell} \leq \ell_i \leq \bar{\ell} < \infty$. Liquidity cannot be raised costlessly at the decision date.*

An aggregate shock can move F , g , ℓ , and K together. The comparative statics below vary one primitive at a time; Section 6 discusses how the data can do the same.

3.2 Rationing by value per dollar

Proposition 1 (Rationing). *Under Assumption 1, there is a multiplier $\lambda \geq 0$ such that an optimal allocation funds every opportunity with $g_i > \lambda \ell_i$, funds none with $g_i < \lambda \ell_i$, and*

randomizes on equality. The value function satisfies

$$V(K) = \min_{\lambda \geq 0} \left\{ \lambda K + \int [g_i - \lambda \ell_i]_+ dF(i) \right\}, \quad (2)$$

is nondecreasing and concave, and $V'(K) = \lambda$ wherever it is differentiable. Opportunities with $0 < g_i < \lambda \ell_i$ are profitable but unfunded. The funding threshold λ is nonincreasing in K , and it is zero once K covers every positive-value opportunity.

The multiplier λ is the shadow value of a dollar at the decision date. An action worth g_i is forgone when the cash it uses would earn more elsewhere in the portfolio. With heterogeneous cash needs, total value is the wrong ranking: a high-balance or high-coupon loan can be passed over precisely because it absorbs more scarce cash. The structure is a continuous knapsack, the same structure that governs Ginnie Mae buyout exercise in Bandyopadhyay (2026d). The contribution here is what it implies once information, verification, and borrower welfare enter.

3.3 Selection and the appearance of success

To connect rationing to information, let $q_i \in [0, 1]$ be the intermediary's posterior probability that the intervention succeeds, for instance that the loan reperforms, and specialize to

$$g(q) = Rq - c, \quad \ell(q) = k - \omega q, \quad R > c \geq 0, \quad k > \omega \geq 0, \quad Rk > \omega c. \quad (3)$$

Success raises the private payoff by R at cost c ; the parameter ω captures offsets to the cash requirement that are larger for better loans (for example, a higher advance rate on a loan expected to reperform). Setting $\omega = 0$ gives a constant cash need. With slight abuse of notation, F now denotes the distribution of q , with density f .

Proposition 2 (Selection). *The optimal policy funds $q > q^*(\lambda) = (c + \lambda k)/(R + \lambda \omega)$, with $dq^*/d\lambda = (Rk - \omega c)/(R + \lambda \omega)^2 > 0$. If the budget binds and q has a positive density near*

the cutoff, a reduction in K lowers the funded share and raises $\mathbb{E}[q \mid x = 1]$. If q is the true success probability under intervention, the success rate among funded interventions rises as access contracts.

Proposition 2 is a selection result. It says nothing about the effect of intervention relative to continuation, which requires both potential outcomes. Its empirical content is that a contraction in liquidity improves the measured performance of the relief that is delivered: cure rates among modified or repurchased loans rise just as fewer borrowers are helped. Performance metrics computed on treated loans therefore cannot be read as evidence that a servicer’s relief improved. Figure 1 and Table 1 illustrate with a uniform posterior.

3.4 Value shocks and liquidity shocks

The rationing structure yields a diagnostic for whether the budget binds. Let the cash need be constant, $\ell_i = k$, so the fundable share is $m = K/k$.

Proposition 3 (Value shocks versus liquidity shocks). *Let $\bar{m} = \Pr(g > 0)$ be the share of positive-value opportunities.*

- (a) *If the budget binds ($m < \bar{m}$), the share of opportunities funded equals m . It rises by $1/k$ per dollar of liquidity and is unchanged by any shock to values that leaves at least a share m of opportunities with positive value.*
- (b) *If the budget is slack ($m > \bar{m}$), the share funded equals $\bar{m}$. It does not respond to liquidity and moves one-for-one with the share of positive-value opportunities as long as that share remains below m .*
- (c) *With heterogeneous cash needs, the cash deployed under a binding budget equals K for every value shock that leaves positive-value cash demand above K .*

Proposition 3 is a double dissociation: liquidity moves constrained intermediaries, and value moves unconstrained ones. A common shock that moves every opportunity’s value in the same direction cannot move two intermediaries’ relief in opposite directions. A common

drain on cash combined with a common rise in value can: the rise in value raises relief where funding is slack, and the drain lowers it where funding binds. The 2020 policy environment supplies both kinds of variation. Pass-through assistance and facility conditions shifted liquidity, while the June 2020 redelivery restriction shifted value for a dated subset of loans and the 2022 rise in mortgage rates lowered the value of repurchase for most of them.

3.5 Where scarcity comes from

Write liquidity as $K = A - M$, where A is usable cash and committed borrowing capacity and M is the cash required by existing obligations over the decision interval, chiefly advances on delinquent loans. A forbearance wave raises M for every issuer. For a depository, A is large and does not contract with delinquency. For a nonbank, A is capped by facility limits that tighten as collateral values fall. The same event that raises the value of repurchase, a surge in delinquencies, also drains the resource needed to exercise it. This is why a regression of relief on contemporaneous delinquency does not isolate the liquidity channel: delinquency is at once a cash demand, a state variable for asset quality, and a determinant of the value of action.

4 Liquidity and the Value of Information

4.1 Setup

Fix a cell of opportunities that are identical on hard information and have success probability p under intervention. Before acting, the intermediary observes a signal with posterior $Q = \Pr(\text{success} \mid \text{signal})$ distributed according to G on $[0, 1]$; Bayes plausibility requires $\mathbb{E}[Q] = p$. Each intervention yields $RQ - c$ and requires cash k , so the fundable share is $m = K/k \in [0, 1]$. I assume $Rp > c$, so that intervention is profitable on hard information alone and the intermediary would act without the signal, and $\Pr(Q = p) = 0$, so that every

signal is either good or bad news.¹ Write $G^{-1}(u) = \inf\{q : G(q) \geq u\}$ for the quantile function and

$$T_G(m) = \int_{1-m}^1 G^{-1}(u) du \quad (4)$$

for the expected number of successes among the top-ranked share m . By Proposition 1, the intermediary funds the highest posteriors first and stops at zero value, so its value with and without the signal is

$$V_G(m) = \int_{1-m}^1 [RG^{-1}(u) - c]_+ du, \quad V_0(m) = m(Rp - c). \quad (5)$$

The allocation value of the signal is $\Delta V(m) = V_G(m) - V_0(m)$.

4.2 The main result

Proposition 4 (Liquidity and the value of information). *Let $m^* = \Pr(Q > p)$. The allocation value of the signal has the following properties.*

- (i) *Decomposition. $\Delta V(m) = \mathcal{A}(m) + \mathcal{S}(m)$, where the ranking value $\mathcal{A}(m) = R[T_G(m) - mp]$ and the screening value $\mathcal{S}(m) = \int_{1-m}^1 [c - RG^{-1}(u)]_+ du$ are nonnegative. $\mathcal{A}$ is concave with $\mathcal{A}(0) = \mathcal{A}(1) = 0$. $\mathcal{S}$ is convex, nondecreasing, and zero whenever $m \leq \Pr(RQ \geq c)$.*
- (ii) *Shape. ΔV is concave, with $\Delta V(0) = 0$ and $\Delta V(1) = \mathbb{E}[c - RQ]_+ \geq 0$. It is nondecreasing on $[0, m^*]$ and nonincreasing on $[m^*, 1]$.*
- (iii) *Complementarity. The marginal value of liquidity is higher with the signal than without it if $K < km^*$ and lower if $K > km^*$. Information and liquidity are complements when liquidity is scarce and substitutes when it is abundant.*
- (iv) *Peak. The maximal allocation value is $\Delta V(m^*) = R\mathbb{E}[Q - p]_+ = \frac{R}{2}\mathbb{E}|Q - p|$, and it consists entirely of ranking value.*

¹If $\Pr(Q = p) > 0$, parts (ii) and (iv) of Proposition 4 hold with m^* replaced by any point of $[\Pr(Q > p), \Pr(Q \geq p)]$, an interval on which ΔV is constant and the marginal values of liquidity with and without the signal are equal. Where V has kinks, marginal values refer to one-sided derivatives.

Three implications follow. First, the value of the same information differs sharply across balance sheets. An intermediary that can fund almost nothing gains little from knowing whom it would help, and one that can fund every opportunity uses information only to decline. The value is largest in between, at the fundable share equal to the share of borrowers whose signal is good news. Second, the statistic that governs the maximal allocation value is the mean absolute deviation of the posterior. It is not the variance, the coefficient of determination, or classification accuracy, and it measures how far the signal moves beliefs, weighted equally in both directions. Third, the regime determines the kind of value information delivers. When liquidity is scarce, information is used to decide whom to help first. As liquidity approaches full funding, it is used increasingly to decide whom not to help, and its value becomes the option to decline actions that are unprofitable on the finer information; if no such actions exist, the value of information vanishes at full funding. Figure 2 illustrates with a continuous posterior.

Corollary 1 (Binary signal). *Let a binary signal have symmetric accuracy $\eta \in (1/2, 1)$, meaning that it is good with probability η after success and bad with probability η after failure, so that $Q \in \{q_L, q_H\}$ with $\Pr(Q = q_H) = \pi_H = p\eta + (1 - p)(1 - \eta)$, and suppose $Rq_L > c$. Then $\Delta V(m) = R(q_H - p)m$ for $m \leq \pi_H$ and $\Delta V(m) = R(p - q_L)(1 - m)$ for $m \geq \pi_H$.*

With $p = 1/2$, $\eta = 0.8$, $R = 1$, and $c = 0.1$, the posteriors are 0.8 and 0.2 and the peak value is 0.15 at $m^* = 1/2$. Concavity implies that if adopting the signal carries a fixed cost ψ paid from a separate operating budget, the intermediary adopts it on an interval of liquidity levels containing m^* whenever $\psi < \Delta V(m^*)$. In the example, a cost of 0.06 makes adoption worthwhile only for fundable shares between 0.2 and 0.8. If the cost is paid from the same scarce cash, the comparison becomes $V_G(m - \psi/k)$ against $V_0(m)$, and adoption is further deterred at low liquidity. Both intermediaries with ample liquidity and intermediaries with very little have weak incentives to invest in borrower information.

4.3 Reliance on information

Proposition 4 concerns the value of information. A closely related object governs how strongly decisions respond to information, which is what loan-level data reveal. Let $r = h(Q)$ be any ranking variable that is nondecreasing in the posterior, such as a value proxy, with $0 < \text{Var}(r) < \infty$. Define the level response $\tau_r(m) = \text{Cov}(x, r)/\text{Var}(r)$, the slope of a linear probability regression of intervention on r , and the relative response $\tau_r(m)/\Pr(x = 1)$.

Proposition 5 (Reliance on information). *Suppose $RQ \geq c$ almost surely, so the intermediary funds the top-ranked share m .*

- (i) *The level response is $\tau_r(m) = \left[\int_{1-m}^1 h(G^{-1}(u)) du - m \mathbb{E}[r] \right] / \text{Var}(r)$. It is concave in m , zero at $m = 0$ and $m = 1$, and maximized where the funded share equals $\Pr(r > \mathbb{E}[r])$. For $r = Q$, $\tau_Q(m) = \mathcal{A}(m)/(R \text{Var}(Q))$: the level response is proportional to the ranking value of information and peaks at m^* .*
- (ii) *The relative response is $\tau_r(m)/m = (\mathbb{E}[r \mid x = 1] - \mathbb{E}[r]) / \text{Var}(r)$, which is nonincreasing in m . As liquidity contracts, the response relative to the base rate weakly rises.*
- (iii) *The difference in intervention rates between borrowers with good and bad news, $D(m) = \Pr(x = 1 \mid Q > p) - \Pr(x = 1 \mid Q < p)$, equals m/m^* for $m \leq m^*$ and $(1 - m)/(1 - m^*)$ for $m \geq m^*$.*

The proof uses only that the intermediary funds a top share of its ranking and that r is monotone in it. The proposition therefore applies to any ranking index, including hard characteristics such as a loan's coupon that enter the ranking through the value of action, with Q replaced by the index. A constrained intermediary keeps ranking and funds less of its ranking. When liquidity contracts below the peak, the level sensitivity of relief to a ranking variable falls toward zero, while the sensitivity relative to the collapsed base rate rises. A flattening of the level relationship between relief and a value proxy is therefore what rationing predicts, not evidence that the intermediary stopped discriminating. It is also not, by itself, a signature of a funding constraint. A value shock that compresses values across

the ranking variable can produce the same flattening, and a uniform shift in values leaves an unconstrained intermediary funding the top of the same ranking (Internet Appendix IA.1.2). The slopes are diagnostic in combination with the dated shocks of Proposition 3.

4.4 Which model is worth more?

Servicers rank delinquent borrowers with statistical and machine-learning models, and the standard measure of such a model's quality is its AUC, the probability that it ranks a randomly chosen success above a randomly chosen failure. Proposition 4 implies that the relevant measure depends on the balance sheet. Define the cumulative accuracy profile $\Gamma_G(m) = T_G(m)/p$, the share of all eventual successes captured by the top-ranked share m . Then the ranking value in Proposition 4 is

$$\mathcal{A}_G(m) = Rp[\Gamma_G(m) - m], \quad (6)$$

proportional to the vertical distance between the accuracy profile and the diagonal at the fundable share.

Proposition 6 (Evaluating scores for a constrained intermediary). *Consider two scores that produce calibrated posteriors $Q_A \sim G_A$ and $Q_B \sim G_B$ with the same mean p .*

- (i) *Score A is worth at least as much as score B at every fundable share $m \in [0, 1]$ and every cost $c \geq 0$ (a comparison that does not involve the no-signal benchmark) if and only if $Q_A \geq_{\text{cx}} Q_B$, that is, if and only if $\Gamma_A(m) \geq \Gamma_B(m)$ for all m . For a binary outcome, this is Blackwell's ordering of the two experiments.*
- (ii) *If $RQ \geq c$ almost surely under a score, its allocation value averaged over fundable shares drawn uniformly from $[0, 1]$ is*

$$\int_0^1 \Delta V_G(m) dm = Rp(1 - p) \left(\text{AUC}_G - \frac{1}{2} \right), \quad (7)$$

where AUC_G counts ties as one-half.

(iii) Under the condition in (ii), if the accuracy profiles cross, there are fundable shares at which each score is strictly more valuable, and AUC ranks the scores by their average value in (ii).

Part (i) says that a score is better for every intermediary only if it is more informative in Blackwell's sense. That is a partial order, and most pairs of models that practitioners compare are not ranked by it. Part (ii) identifies what AUC measures: the value of a score to an intermediary whose liquidity is equally likely to take any value. Hernández-Orallo, Flach, and Ferri (2012) show that AUC is linear in expected loss when the classification threshold is chosen by a uniformly distributed rate. Equation (7) is the economic version of that result for a capacity-constrained intermediary, with the payoff scale and the base rate made explicit. A servicer that knows its liquidity should evaluate models at that liquidity.

Table 2 and Figure 3 illustrate with two calibrated scores with the same mean $p = 0.5$. Score A is precise at the top: it identifies 20% of borrowers who succeed with certainty and assigns everyone else 0.375. Score B is broad: it splits borrowers into halves with posteriors 0.8 and 0.2. Score B has the higher AUC, 0.80 against 0.70. At a fundable share of 7.2%, the nonbank exercise rate in June 2020 in the sample of Section 6, score A is worth two-thirds more than score B . At 86.2%, the depository rate, score B is worth 2.4 times as much. The two issuer types would choose different models, and both would be right.

4.5 When information expands funding

So far, information allocates a fixed budget. It can also change the budget. Suppose each funded loan can be pledged to a competitive lender against a claim that pays L if the intervention succeeds, with $L < R$ and $L < k$, so that each action requires some internal cash. Internal liquidity is $K_0 \in (0, k - Lp)$, and $RQ > c$ almost surely, so the budget binds.

Proposition 7 (When information is financeable). *(a) If the lender prices the facility on*

hard information, advancing Lp per funded loan, then whatever the intermediary's signal, the fundable share is $m_H = K_0/(k - Lp)$.

(b) If the lender knows the intermediary's information technology G and observes which loans are funded, but not their posteriors, the competitive advance on a pool of the top-ranked share m is $LT_G(m)/m$ per loan. The fundable share m_R solves $km - LT_G(m) = K_0$, and loan-level disclosure of posteriors yields the same m_R .

(c) $m_R \geq m_H$, strictly if Q is nondegenerate.

(d) The intermediary's private value is $V_G(m_R)$ under (b) and $V_G(m_H) - L[T_G(m_H) - m_{HP}]$ under (a), and the former is at least as large.

Information that the lender does not recognize is valuable for allocation but not for financing. Under (a), the intermediary also cross-subsidizes the lender, because the funded pool is better than the facility's price assumes. Under (b), the lender infers the pool's quality from the intermediary's known selection rule, and the budget expands although no borrower-level information changes hands. What must be verifiable is the technology and its use, not each score. The distinction matters for how intermediaries adopt predictive technology. An internal model that facility lenders cannot validate improves who is helped but not how many can be helped. A model whose existence and use lenders can verify, through validation, audit, or shared data standards, can do both. In the numerical example of Corollary 1, with $k = 1$, $L = 0.6$, and $K_0 = 0.2$, recognition raises the fundable share from 28.6% to 38.5%. The assumption in (b) that the lender observes which loans are funded is natural for buyout facilities, where the funded loans are the collateral. If the intermediary instead chooses which of its loans to pledge, pooled pricing invites adverse selection in the pledged set (Myers and Majluf, 1984; DeMarzo and Duffie, 1999), and loan-level verification regains a role.

5 Borrowers, Stakeholders, and Policy

5.1 Private selection and borrower surplus

Let b_i be the borrower's incremental surplus from intervention, v_i the gain to other private claimants, and t_i the net gain to taxpayers, all relative to the same continuation. Total incremental surplus is

$$w_i = g_i + b_i + v_i + t_i, \tag{8}$$

with real resource costs counted once and transfers netting out. A planner facing the same funding technology ranks by w_i/ℓ_i ; the intermediary ranks by g_i/ℓ_i . In the information setting of Section 4, write $\beta(q) = \mathbb{E}[b \mid Q = q]$ and $\bar{w}(q) = \mathbb{E}[w \mid Q = q]$.

Proposition 8 (Private selection and borrower surplus). *Suppose $RQ \geq c$ almost surely, so the intermediary funds the top-ranked share m .*

- (i) *If β is nonincreasing, the signal weakly lowers borrower surplus at every fundable share, strictly for $m \in (0, 1)$ if $\beta(Q)$ is nondegenerate.*
- (ii) *If $\bar{w}$ is nonincreasing, the signal weakly lowers total surplus at every fundable share while raising private value.*
- (iii) *A planner with the same signal and the same budget weakly raises total surplus relative to allocation without the signal.*

The condition in (i) is the self-cure logic of Adelino, Gerardi, and Willen (2013). Borrowers with the best prospects under intervention are often those most likely to recover without it, so their incremental gain is smallest. When that ordering holds, better private selection systematically steers relief away from the borrowers who gain most. For example, let half the loans have posterior 0.8 and half 0.2, with $R = 1$, $c = 0.1$, borrower gains of 0.1 and 0.9, and capacity for half the loans. Without the signal, private value is 0.20 and borrower surplus 0.25 per loan. With it, private value is 0.35 and borrower surplus 0.05; total surplus falls from 0.45 to 0.40. Part (iii) locates the problem. The welfare loss arises from using information to

maximize private recovery value, not from the information itself, so improvements in servicing technology should be evaluated with borrower outcomes rather than with investor returns or the performance of treated loans.

5.2 Displacement within the intermediary

Under a binding budget, relief to one loan is funded by relief denied to another. With a constant cash need, the number of funded interventions is fixed locally, and better information about one group of borrowers changes who is helped without changing how many. With heterogeneous cash needs, the budget conserves dollars rather than counts:

$$\sum_{\text{groups } h} d \left[\int_{i \in h} \ell_i x_i dF(i) \right] = 0. \quad (9)$$

Treatment of one loan therefore affects others serviced by the same constrained intermediary. A loan-level estimate that ignores this interference can overstate the effect of an institution-wide change, such as the adoption of a new model or the arrival of new funding, and a credible evaluation must measure effects across the intermediary's whole book.

5.3 Liquidity support versus ex post compensation

Policy can support relief with liquidity, which supplies cash at the decision date, or with compensation paid after an action is completed, such as the incentive payments per completed modification that servicers received under the Home Affordable Modification Program (Agarwal et al., 2017). The two are often treated as substitutes. They are not when funding binds.

Proposition 9 (Liquidity versus compensation). *Consider an ex post payment s per completed action and a liquidity injection ΔK .*

- (a) Binding budget, constant cash need. *Suppose values have a positive density at the margin. The payment leaves the set of funded actions unchanged and raises the shadow*

value of cash by s/k . The injection raises the funded share by $\Delta K/k$ as long as the budget still binds.

- (b) Binding budget, heterogeneous cash needs. A flat payment leaves the cash deployed equal to K and weakly raises the number of actions only by tilting relief toward actions with lower cash needs. A payment proportional to the cash need leaves the allocation unchanged.
- (c) Elastic funding. Suppose external funds B cost $\Phi(B)$ with Φ increasing and strictly convex, cash needs are constant, values are $g = RQ - c$ with Q continuously distributed, and at the optimum the intermediary borrows and funds an interior share. For every payment $s > 0$ there is an injection that produces the same funded share. The payment raises the intermediary's external borrowing and the injection lowers it.

Under a hard budget, compensation paid after completion is inframarginal: it raises the value of each action but not the cash available to take one, so the budget spends it on the same actions or on cheaper ones. Liquidity support expands relief directly. When the intermediary can borrow at a rising marginal cost, compensation does expand relief, because it raises the value of actions financed at the margin. The volume of relief then cannot distinguish the instruments, but the intermediary's facility draws can. A payment made at the time of the action combines the two: it is both liquidity and compensation. The Pass-Through Assistance Program was liquidity priced as a last resort, which places it in case (c) with a steep Φ . At a continuous interior cutoff, the social return to a marginal dollar of liquidity is $\bar{w}(q^*)/\ell(q^*)$. That is a local expression; a full benefit-cost calculation also requires fiscal costs, default risk, and the counterfactual private financing.

5.4 Menus and staff capacity

Relief involves a menu. Let $a \in \mathcal{A}_i$ index continuation, repurchase, forbearance, modification, sale, and liquidation, as available for loan i . Each action has private value g_{ia} , cash need ℓ_{ia} , and staff requirement d_{ia} . With aggregate cash K and staff capacity D , the optimal action

solves

$$a_i^* \in \arg \max_{a \in \mathcal{A}_i} \{g_{ia} - \lambda \ell_{ia} - \mu d_{ia}\}, \quad (10)$$

where $\lambda, \mu \geq 0$ are the two shadow prices. A decline in completed relief can reflect either one, and Agarwal et al. (2017) show that servicers' organizational capacity limited modifications under HAMP. The two constraints are observationally equivalent when staff and cash requirements are proportional across actions. If $d_{ia} = \varrho \ell_{ia}$ for all actions, choices depend only on $\lambda + \varrho \mu$. Separating them requires independent variation in cash and processing intensity or a direct measure of one constraint, which is why documentation and queue data are part of the mechanism rather than controls (Bandyopadhyay, 2026b).

6 Evidence

This section takes the rationing, reliance, and value-shock predictions to Ginnie Mae buyout decisions from 2019 to 2022, a period with a liquidity shock concentrated on nonbank issuers in 2020, its reversal in 2021, and two shocks to the value of repurchase, Ginnie Mae's June 2020 restriction on redelivery and the rise in mortgage rates in 2022. Before turning to the data, I state which statistic each kind of shock should move. Under a liquidity shock, the funded share moves only at constrained intermediaries, and the ranking of options stays put (Propositions 1 and 3). Under a value shock, the funded share moves at unconstrained intermediaries, and the ranking itself can change for everyone; if the shock pushes part of the opportunity set below zero, intermediaries that fund every profitable action start to use their ranking to decline (Propositions 3 and 4). Section 6.1 describes the data and a ranking index. Section 6.2 follows the ranking of constrained issuers through the contraction and recovery of their budgets. Section 6.3 tests the rule's other half, that scarce cash tilts the ranking against the actions that use most of it, and Section 6.4 relates exercise to each issuer's own cash drain. Section 6.5 turns to unconstrained issuers and to the value shocks. Section 6.6 shows that the options constrained issuers postponed in 2020 were exercised later

according to the same ranking, and Section 6.7 documents repurchases that unconstrained issuers made and constrained issuers did not. Section 6.8 connects the results to companion work, Section 6.9 puts the estimates together in a quantitative illustration, and Section 6.10 collects the predictions and the tests that remain.

6.1 Data and a ranking index

The data are Ginnie Mae’s public monthly loan-level disclosures for single-family pools from December 2013 to December 2022, in files assembled from the published disclosures for Bandyopadhyay, Kim, and Smith (2026). I use them in two forms. The first has one observation per buyout option, in the month the option vests because the loan first reaches three missed payments, with the loan’s hard characteristics at that date, the month and reason of any removal from its pool over the following twelve months, and its last recorded status through December 2022. The second records, for every month, the stock of loans that are at least three payments delinquent and therefore eligible for repurchase. A repurchase is a removal under reason 2 before any cure. The main outcome is exercise within three months of vesting, which captures the decision at vesting and is observed for every cohort through September 2022; Section 6.6 studies later exercise. The sample contains the 1,922,668 options that vested from January 2019 to September 2022 at the 120 issuers classified as depositories or nonbanks: 1,298,771 at nonbanks and 623,897 at depositories. The classification follows Bandyopadhyay (2026d), extended by charter to the largest issuers it does not cover (Internet Appendix IA.4.1). Nonbanks fund advances and repurchases from capped facilities; depositories fund them from deposits (Kim et al., 2018; Jiang et al., 2020); Section 6.4 adds a direct measure of each issuer’s cash drain.

The model ranks actions by value per dollar of cash, which is not observed. Proposition 1 says how to recover the ranking: a constrained intermediary funds the top of its ranking, so which options it exercises within a decision cell reveals its ranking whatever the size of its budget. I therefore estimate the ranking from nonbanks’ own earlier decisions. The ranking

index is the linear predictor of a logit of exercise on hard loan characteristics at vesting, estimated with issuer-month fixed effects on the 298,723 nonbank options that vested from 2016 to 2018 in issuer-months in which some but not all options were exercised. The fixed effects absorb each issuer-month’s budget, so only the ranking is learned. The inputs are the loan’s note rate minus the 30-year mortgage rate in the vesting month, which governs the premium at which a repurchased loan can be redelivered, the change in the ten-year Treasury rate since origination, loan age, balance, credit score, current loan-to-value ratio, judicial-foreclosure state, insurance program, pool issue type, first-time-buyer status, loan purpose, prior missed payments, and the annual insurance premium. The mortgage spread dominates: a one-point higher spread raises the log-odds of exercise by 1.24 within an issuer-month, as the redelivery value of a repurchased loan implies (Bandyopadhyay, Kim, and Smith, 2026), and VA loans rank lower. On decisions after the estimation window, the index’s size-weighted area under the ROC curve within issuer-month cells is 0.73 for nonbank decisions in 2019, 0.68 in 2020, 0.67 in 2021, and 0.78 in 2022. For depositories, which exercised four in five options and selected little before 2022, it is 0.54 in 2019. The index is blind to forbearance, which does not occur in the training period, and a repurchase of a loan in forbearance ties up cash for longer (Internet Appendix IA.2). I therefore define decision cells as issuer by month by forbearance status, so that the ranking is compared only among options that use cash in the same way. The disclosures flag forbearance from May 2020; in March and April 2020 the cells are issuer-months. Issuers also rank on information the disclosures do not contain (Bandyopadhyay, Kim, and Smith, 2026); the estimation assumes that this part of the ranking is logistic, independent of the index, and stable. I standardize the index over all options in the sample.

6.2 Rationing along a fixed ranking

Propositions 1 and 5 imply that when a constrained intermediary’s budget shrinks, it keeps ranking opportunities the same way and funds a smaller top share of its ranking, and that it

moves back down the same ranking when the budget recovers. Four alternatives are natural. Under random rationing, a constrained issuer funds a random subset of the options it would otherwise exercise, for instance in the order in which they are processed; its response to value then falls in proportion to its funded share, and its response relative to the base rate does not rise. If an issuer in distress stops discriminating, both responses fall. If the shock changes which loans are valuable, the weight on the index changes. Finally, a uniform decline in the value of every option would lead even an unconstrained issuer to fund a smaller top share of the same ranking, and it produces the same pattern as rationing (Internet Appendix IA.1.2). The comparisons below bear on the first three alternatives; Section 6.3 adds a prediction that none of the four makes; and the fourth is addressed by the contrast with depositories and by the value shocks in Section 6.5.

The fixed-ranking prediction is implemented in two steps. Let x_{ijt} indicate exercise of option i in decision cell j (an issuer, month, and forbearance status) and let r_{ijt} be its ranking index. First, I estimate a logit with cell fixed effects,

$$\Pr(x_{ijt} = 1) = \Lambda(\alpha_{jt} + \beta r_{ijt}), \quad (11)$$

on nonbank decisions from March 2019 to February 2020. The fixed effects absorb each cell's budget, so β measures how heavily the issuers' ranking weights the index; the logistic error represents the part of their ranking that the index misses. Second, holding β at its pre-period value, for every later cell I choose α_{jt} to match the observed funded share and compute the implied level response of exercise to the index, the within-cell covariance of exercise with the index divided by its variance, and the implied relative response, the level response divided by the funded share. By Proposition 5(ii), the relative response is also the selection statistic of Proposition 2: the mean index of exercised options relative to all options, scaled by the variance of the index. The predictions use the pre-period weight, each cell's funded share, and each cell's distribution of the index, but not which loans were exercised after February

2020. For the random-rationing benchmark, each issuer funds, along the same ranking, the options it would have funded at its own pre-period funded share, and then exercises a random fraction of them that matches its observed funded share.

Figure 4 and Table 3 report the results. Nonbanks' funded share fell from 40.0% in the year before the pandemic to 8.3% in March through June 2020. Their level response to the index fell from 0.201 to 0.043, as every alternative predicts. The response relative to the base rate, which separates the alternatives, was 0.50 before the shock and 0.52 in the freeze. Rationing along the pre-shock ranking, given each cell's funded share, predicts a rise to 0.75; random rationing predicts a fall to 0.40; a loss of discrimination predicts a larger fall. The actual response lies between the first two: the difference from the fixed-ranking prediction is -0.23 (bootstrap s.e. 0.25) and from the random-rationing benchmark 0.11 (0.16), and the latent weight on the index changed by -0.49 (s.e. 0.31) against a pre-period weight of 1.45 (0.29). The freeze months are therefore consistent with rationing along a ranking that the index captures less well than before the shock, or with a mixture of ranked and random rationing, and they do not by themselves distinguish the two. The point estimate shows no fall in the relative response; the bootstrap does not exclude a fall of up to about 0.17.

The aggregate conceals heterogeneity across issuers. Internet Appendix Table IA.3 reports the comparison separately for the five nonbanks with the most options before the shock, among issuers that exercised at least 5% of them (an issuer that exercised almost none has no estimable pre-shock weight), and for the other 78 nonbanks as a group. The five held 59% of pre-shock options and 67% of freeze options, and each row re-estimates the pre-shock weight on that issuer's own decisions. Three of the five, and the group of small issuers, show a freeze response at or above the fixed-ranking prediction. For the issuer with the most pre-shock options, whose funded share fell from 53.4% to 4.2%, the response rose from 0.69 to 1.53, between the fixed-ranking prediction of 2.06 and the random-rationing benchmark of 0.74 and closer to the former in the change from its pre-shock level. The exception is the second-largest issuer, whose funded share fell to 2.4% and whose response fell to 0.38,

below both benchmarks (2.11 and 0.61); it accounts for the aggregate’s shortfall from the fixed ranking. Month by month, the nonbank response was below both benchmarks in March and April, when the cells cannot separate loans in forbearance from the rest, and moved toward the fixed ranking in May and June (0.41 against 0.54 and 0.37; 0.61 against 0.88 and 0.39). Internet Appendix Table IA.2 varies the ranking variable, the decision cell, the exercise window, and the sample, and Internet Appendix Table IA.5 repeats the test on the monthly stock of eligible loans.

The cycle then reversed. In the second half of 2020 the nonbank funded share rose to 15.9% and the relative response to 0.67, close to the fixed-ranking prediction of 0.71 (a difference of -0.03 , s.e. 0.11) and 0.20 above the random-rationing benchmark (s.e. 0.08; 95% percentile interval 0.00 to 0.31). This is the period in which the comparison is most informative, with the two benchmarks well separated and the narrowest bootstrap interval, and it favors the fixed ranking. In the first half of 2021 the funded share recovered to 32.4% and the response fell to 0.43; the fixed ranking, with the pre-shock weight, each cell’s funded share, and no information on which loans were exercised after February 2020, predicts 0.45, a difference of -0.02 (s.e. 0.05). In the second half of 2021 the response was 0.37 against a prediction of 0.40. Once funded shares are back near their pre-shock levels the two benchmarks nearly coincide, so 2021 is consistent with both; what it shows is that the response moved back to where the pre-shock weight and the recovered funded share put it. The latent weight on the index changed by -0.09 (s.e. 0.24), -0.07 (0.21), and -0.19 (0.13) in these three periods. The fixed-ranking computation also holds each period’s distribution of the index fixed; within-cell dispersion of the index was similar before the shock and during the freeze (0.38 and 0.37) and larger in 2021 (0.52 and 0.60), so the predictions in Table 3 should be read as reflecting the funded share and the index distribution together.

6.3 The price of cash in the ranking

Proposition 1 ranks actions by value per dollar of cash, not by value. With heterogeneous cash needs, the cutoff $g_i > \lambda \ell_i$ tilts against actions that absorb more cash as the shadow price λ rises, so a loan of a given value should become less likely to be repurchased, relative to a cheaper loan of the same value, exactly when the constrained issuer’s cash is scarcest. The cash a repurchase requires is the loan’s unpaid balance. Random rationing, a loss of discrimination, and a value shock make no such prediction: none of them changes the weight on balance holding value fixed, and depositories, whose λ is zero, should show no tilt.

Table 4 adds log balance to the cell-fixed-effects logit of equation (11) with period interactions for both regressors. Before the shock, nonbanks’ weight on balance was 0.19 (s.e. 0.11): holding the index fixed, larger loans were if anything slightly more likely to be repurchased. In the freeze the weight fell by 0.32 (s.e. 0.06), five standard errors from zero and the sharpest of the changes the paper reports, and it returned to its pre-shock level in the release (a change of 0.07, s.e. 0.06) and stayed there through 2021. Depositories’ weight on balance was -0.01 (0.01) before the shock and changed by -0.02 (0.03) in the freeze. The tilt is not an artifact of forbearance. In March and April the cells cannot separate loans in forbearance from the rest, and take-up may have varied with balance; in May and June, when forbearance status enters the cell definition, the change in nonbanks’ weight was -0.32 (s.e. 0.08), against -0.31 (0.15) in March and April, and depositories’ was -0.02 (0.04). The constrained issuers did not stop ranking; they raised the price of cash in their ranking for four months and then lowered it again. This is the signature of a binding budget that none of the alternatives in Section 6.2 produces, and it rests less on the index than the comparisons of that section do: balance enters the model directly as the cash requirement, and the test asks only whether its weight, conditional on the index, moved with the budget. A value shock that lowered the value of large loans more than small ones in exactly these four months, and reversed thereafter, could also move the weight conditional on the index; but such a shock would have been common to both types’ loans, and depositories’ weight did not move.

6.4 Funding pressure within issuer types

Issuer type classifies funding conditions but does not measure them. The disclosures allow one measure of the cash that existing obligations absorb, the M of Section 3: the principal and interest an issuer has advanced on its delinquent loans and not yet recovered, relative to the balance of its portfolio. For nonbanks, advances outstanding averaged 0.22% of balance before the pandemic, 0.31% in the freeze, and 0.55% in the first half of 2021; for depositories, 0.06%, 0.17%, and 0.14%. Within issuer and month, a rise in a nonbank’s advances lowers its funded share at vesting by 9.2 percentage points per log point (coefficient -9.2 , s.e. 3.1), or 4.1 points per within-issuer standard deviation of the measure, and its monthly exercise rate on the stock of eligible loans by 4.1 points per log point (s.e. 1.6). The corresponding depository coefficients are -3.2 (3.5) and -1.5 (3.8). The difference between types in the funded-share response, -6.0 (s.e. 4.6), is not statistically significant, and the nonbank estimate is larger in magnitude with the issuer’s serious-delinquency rate as a control (-15.1 , s.e. 5.4). Within issuer, the weight on the index also falls with a nonbank’s advance burden (-0.97 per log point, s.e. 0.47) and not significantly with a depository’s (-0.10 , 0.17), the within-type counterpart of the decline in the latent weight during the freeze. Internet Appendix Table IA.4 reports alternative measures.

These within-type estimates are small next to the 2020 collapse. The within-issuer standard deviation of the log advance burden is 0.44, and it moves the funded share by a few percentage points, whereas the freeze lowered the funded shares of the five largest nonbanks by between twenty and sixty points at once (Internet Appendix Table IA.3). Two cross-sectional exposures to the shock, the share of an issuer’s portfolio in forbearance in May 2020 and the rise in unemployment across the states in which its loans sit, are not significantly related to the size of an issuer’s collapse within type, although the standard errors are wide (Internet Appendix IA.4.6). The shock that bound in 2020 therefore appears to have been largely common to nonbanks as a class, consistent with the accounts of warehouse and servicing-advance facilities in Kim et al. (2018) and Kim et al. (2024), rather than

proportional to each issuer’s own cash drain. That is why the funding channel is identified here from the contrast between types, from the dynamics of the ranking through the cycle, and from the value-per-dollar tilt, and why the interaction test of Section 6.10, which needs issuer-level variation in usable funding, awaits facility data.

6.5 Unconstrained issuers and value shocks

Depositories raised their exercise during the freeze, from 80.5% to 85.3%. For an unconstrained intermediary, the funded share is the share of options with positive value (Proposition 3(b)), so a rising funded share indicates that the value of exercise rose, as falling rates would imply. The inference requires two assumptions: that depositories’ funding did not bind, which the positive size gradient in Section 6.8 suggests may hold only approximately for some of them, and that the change in value was common to both types’ loans, which holds by construction within pools, where the within-pool comparison in Section 6.8 shows the same divergence. A uniform decline in values, the alternative that the nonbank comparisons cannot rule out, would have lowered depositories’ exercise instead. By Proposition 3, a common shock that moved every loan’s value in the same direction cannot move the two types’ exercise in opposite directions, while a common drain on cash that bound for one type and not the other, combined with rising values, does. Depositories’ weight on the index, 0.13 before the shock, fell by 0.35 (s.e. 0.16) in the freeze, to a point estimate below zero; with 85.3% of options exercised, the weight rests on the few options they declined, and their level response stayed within two hundredths of zero from the freeze through 2021.

Two dated shocks to value allow the converse comparison. The first is the redelivery restriction of June 29, 2020, which lowered the value of repurchasing loans that had entered forbearance on or after March 1 once repurchased on or after July 1. In the second half of 2020, depositories’ funded share fell from 85.3% to 79.2%, as Proposition 3(b) implies when the value of part of the opportunity set falls. The restriction targets loans in forbearance, which is also the variable that defines the decision cells, so its effect on the ranking within

cells is not identified separately from the return of funding and the record-low rates of the same months.

The second value shock is the rise in mortgage rates in 2022. The 30-year rate averaged 3.1% in December 2021 and 6.1% in September 2022,² so that 86% of the options that vested at nonbanks in 2022 and 74% of those at depositories were on loans whose note rate was below the current mortgage rate, against about 1% in the freeze. A repurchased loan is valuable largely because it can be redelivered once it reperforms, and a loan with a below-market rate cannot be redelivered at a premium, so much of the eligible stock lost its value. Depositories' funded share fell from 81.0% in the second half of 2021 to 51.9% in 2022, as Proposition 3(b) predicts for an unconstrained intermediary when the share of positive-value options falls. Proposition 4 predicts something sharper. Once part of the opportunity set has negative value, information acquires screening value, and an intermediary that can fund every profitable action starts to use its ranking to decide which actions to decline. Depositories' level response to the index, 0.010 before the pandemic and within a few hundredths of zero through 2021, rose to 0.177 in 2022 (Figure 5, Panel B), and their relative response rose to 0.34 against 0.03 under the pre-shock weight (a difference of 0.31, s.e. 0.12; percentile interval 0.10 to 0.59). The ranking they began to apply is the mortgage spread itself: in 2022 the spread alone ranks their decisions as well as the index (within-cell AUCs in 2022 of 0.79 for the spread and 0.77 for the index at depositories, and 0.79 and 0.78 at nonbanks). Nonbanks' weight on the index also rose at vesting, by 0.60 (s.e. 0.20), as a value shock that reorders opportunities implies for every intermediary; on the monthly stock of eligible loans, whose index is measured at vesting and is stale for loans that vested years earlier, the nonbank weight falls in 2022 (Internet Appendix Table IA.5). Within cells, the index in 2022 is essentially the coupon, so the data cannot distinguish declining negative-value loans from re-ranking on the coupon; both are what a value shock does. In 2020, when the budget moved, depositories' level response to the index stayed near zero; in 2022, when value moved, they began to select, in the way

²Monthly means of Freddie Mac's Primary Mortgage Market Survey, retrieved from FRED (series MORTGAGE30US).

that screening on a ranking implies. Nonbanks' funded share also fell, to 7.2%. For them the value shock cannot be separated from funding costs: an issuer that finances repurchased loans at market rates also faced higher carrying costs in 2022, which lower the private value of repurchase more for nonbanks than for depositories funded by deposits.

6.6 Postponed exercise

The freeze delayed exercise more than it prevented it. Of the options that vested at nonbanks from March to June 2020, 8.3% were exercised within three months of vesting, 30.8% within twelve months, and 40.2% by December 2022. Before the pandemic the corresponding shares were 40.0%, 43.3%, and 51.3% (Figure 5, Panel A). Among the freeze options not exercised at vesting and still in their pools, 26.4% were exercised in the fourth through twelfth months and 37.5% by December 2022, against 5.2% and 21.4% for the corresponding pre-shock options. About 40% of the delayed exercises took place between October and December 2020; almost none took place in July through September.

The model says in what order a backlog is worked down. If the ranking is fixed and the budget expands, the options funded next are the highest-ranked ones not yet funded; under random rationing, the queue also retains top-ranked options, so later exercise still loads on the index; only a loss of discrimination makes later exercise unrelated to it. Table 5 reports, for options not exercised within three months of vesting and still in their pools, a logit of exercise in the following nine months on the index, with cell fixed effects. For nonbank options that vested in the freeze, the weight on the index is 1.05 (s.e. 0.25), about twice the weight in the pre-shock queue, 0.56 (0.15), and about seven-tenths of the pre-shock weight on exercise at vesting, 1.45. At depositories it is 0.39 (0.24) for the freeze cohort. The options that constrained issuers postponed were exercised later with a weight on the ranking close to the one their decisions had carried before the shock; the issuers had not stopped discriminating. Liquidity changed the timing of repurchases more than their composition, as the dynamic version of the model implies when the budget recovers (Internet Appendix IA.2). The delay

was not free: the issuer kept advancing payments on loans it could not repurchase during the months of delay, and at a common horizon of 24 months the freeze cohort’s exercise, 40.1%, still fell 8.4 percentage points short of the pre-shock cohort’s 48.5%. Of the 31.7-point gap between the two cohorts at three months, about 73% had closed within two years.

6.7 Repurchases depositories made and nonbanks did not

Forbearance gives a view of rationing when the budget binds over several periods. The disclosures flag loans in forbearance from May 2020. In May and June, 90% of the options that vested at nonbanks were on loans in forbearance, and 88% at depositories. Depositories exercised 87.6% of them within three months, more than the 72.1% of options on loans not in forbearance. Nonbanks exercised 3.9%, against 38.6% of options on loans not in forbearance (Table 6). Within issuer-months, forbearance is associated with a probability of exercise 27.2 percentage points lower at nonbanks (s.e. 10.3) and 13.2 points higher at depositories (3.0); conditioning also on deciles of the ranking index leaves both differences essentially unchanged. Because the redelivery restriction lowered the value of repurchasing these loans from July 1, the cleaner comparison is exercise by June 30: depositories had repurchased 76.4% of the May–June options on loans in forbearance by then, and nonbanks 1.1%, against 66.8% and 37.1% for options on loans not in forbearance. By December 2022, nonbanks had exercised 38.5% of the options on loans in forbearance and 52.0% of the others.

The repurchases that depositories made most often were the ones nonbanks made least often. Depositories fund repurchased loans on their own balance sheets, so the value of these repurchases to them need not equal their value to a nonbank, and their profitability is not observed. What the contrast shows is that a repurchase attractive enough for unconstrained issuers to make more often than any other was one that constrained issuers left aside, and then made later. The cash required at repurchase is the loan’s balance in either case, so a static budget does not explain the difference. A multi-period budget does. A loan in forbearance cannot be modified and redelivered until its forbearance ends, a constraint one issuer cited in

its filings as its reason not to repurchase (Bandyopadhyay, 2026d). Its repurchase therefore ties up cash longer, and in the dynamic version of the model (Internet Appendix IA.2) an action’s cost in scarce cash grows with how long it holds funds. On that interpretation, the pattern is the multi-period analogue of Proposition 1: actions with positive value and low value per dollar-month of cash are the ones a constrained intermediary leaves unfunded. Differences in the eligibility of loans in forbearance as facility collateral would work in the same direction; the disclosures do not allow the two to be separated.

6.8 Evidence from companion work

Three further findings from Bandyopadhyay (2026d), which studies all 3,138,350 option-exercise decisions from December 2013 to September 2020, fit the same structure. First, within the same pool and month, nonbank-serviced loans were 30.2 percentage points less likely to be repurchased (s.e. 10.5), and the paper’s baseline difference-in-differences estimate is -27.4 percentage points (s.e. 8.5); that paper reports the pre-period trend difference between the two types and a sensitivity analysis for it. Second, the within-type size gradients reverse sign. A larger delinquent book raised a depository’s exercise by 14.4 percentage points per standard deviation of scale and lowered a nonbank’s by 8.0, and the estimated difference between the two gradients is 22.5 points (s.e. 7.3). Under a binding budget, the exercise rate is governed by liquidity per dollar of vesting collateral, which falls with scale for an institution whose facilities are capped and not for one whose funding grows with its size; the model signs the difference between the gradients, not the positive depository gradient itself. Third, the unfunded options did not move to better-capitalized holders. Servicing transfers rose after the shock, but 93% of transfers by nonbanks went to another nonbank, and relative to the same issuers’ 2019 behavior an estimated \$26.8 billion of delinquent principal remained unexercised. Because the option cannot be sold apart from the servicing, the profitable but unfunded set of Proposition 1 stayed with the constrained holder.

The information results also require that intermediaries hold decision-relevant borrower

information and exercise discretion over relief, and both premises are documented. Bandyopadhyay, Kim, and Smith (2026) show that before the pandemic, a one percentage point higher interest rate spread raised the probability of buyout by seven to nine percentage points, and that issuers acquired soft information during delinquent servicing that predicted reperformance and used it in their buyout decisions. Using one servicer’s call-center transcripts for 19,418 borrowers, Bandyopadhyay (2026a) shows that servicer discretion over pandemic forbearance operated through documentation, verification, and referral rather than outright denial. Bandyopadhyay (2026c) follows 12,691 loans at the same servicer and finds that an April 2020 forbearance conversation predicts a 9 to 15 percentage point higher probability of subsequent modification, concentrated among borrowers already close to foreclosure and only where the servicer retained foreclosure discretion. Bandyopadhyay (2026b) shows that documentation frictions in loss-mitigation applications raise foreclosure hazard.

6.9 A quantitative illustration

The two issuer types’ funded shares in the freeze, 8.3% and 85.3%, lie on opposite sides of m^* for any signal whose share of good news lies between them. By Proposition 4, introducing such a signal would raise the marginal value of liquidity for a nonbank and lower it for a depository. Table 2 makes the implication concrete for model choice. The signal structures in the table are hypothetical, and the fundable shares are set near the observed exercise rates only to fix ideas. For a constrained servicer, AUC is not a close proxy for the value of a model; it is an average over balance sheets the servicer does not have. The 2022 episode shows the other side of the same result: when a value shock pushed part of the opportunity set below zero, the unconstrained issuers’ use of information shifted from nothing to screening.

6.10 Predictions and tests

Table 7 collects the model’s predictions, the observable each implies, and the current state of evidence. The first and third are supported by the evidence above, the fourth by the release

and recovery of 2020–2021 rather than by the freeze itself, and the fifth in its screening form. The others remain to be tested, and the rest of this section describes how they can be tested with dated measures of issuer funding, outcomes on both sides of the decision, and, for the borrower-welfare prediction, servicer records.

Risk set and outcomes. The unit of analysis is a loan at the date an action becomes available, such as the vesting of the buyout option. For any funding event, the primary population is a cohort frozen on information dated before the event, so that the population is not conditioned on delinquency caused by the same shock. Outcomes are measured at fixed horizons: delivered modification, sustained performance, completed foreclosure, and redelivery. Loss of observation is recorded as such, and disappearance from a pool is not treated as default or cure. Internet Appendix IA.3 gives the construction protocol.

Value and cash. Empirical analogues of g and ℓ use only information available at the decision: the difference in expected discounted cash flows between action and continuation, and current cash required net of immediately drawable financing. Expected recovery, holding period, funding rates, and the redelivery regime in force at the decision date enter separately. A risk score estimated on historical outcomes, $\hat{q}_{it}^{\text{obs}} = \widehat{\text{Pr}}(Y_{i,t+h} = 1 \mid H_{it}, S_{it})$, predicts outcomes under the historical policy. It is neither the success probability under intervention nor the treatment effect without further assumptions, and the tests below use it only as a ranking variable. Text or other soft signals must be time-stamped before the decision.

Funding variation and the information interaction. Propositions 4 and 5 imply that a funding shock changes the use of prior information with a sign that depends on the regime. With issuer funding variation Z_{jt} and a prior signal $S_{i,t-1}$,

$$\begin{aligned}
y_{ijt} = & \alpha_j + \phi_{c(i),t} + \pi Z_{jt} + \gamma S_{i,t-1} + \theta Z_{jt} S_{i,t-1} + \theta_2 Z_{jt} S_{i,t-1} \mathbf{1}\{m_{j0} > \hat{m}^*\} \\
& + f(H_{i,t-1}, \hat{g}_{i,t-1}, \hat{\ell}_{i,t-1}) + \varepsilon_{ijt},
\end{aligned} \tag{12}$$

where y_{ijt} indicates that issuer j acted on loan i in month t , $\phi_{c(i),t}$ are risk-cell-by-month effects, m_{j0} is the issuer’s pre-event funded share, and $\hat{m}^*$ is the share of loans in the cell whose prior signal exceeds the cell mean, the switch point in Proposition 5(i). The model predicts $\theta > 0$ and $\theta + \theta_2 < 0$: more funding raises reliance on information for issuers below the switch point and lowers it above. The prediction is a sign reversal at a location fixed in advance, which is harder to produce by specification search than a single interaction. Credible Z_{jt} must shift usable funding without shifting loan values. Candidates include disruptions at an issuer’s pre-existing facility lenders and mechanically determined facility limits, and each requires a verified first stage and an exclusion argument. The dated value shocks of Section 6.5 supply the complementary comparison that does not require an excluded funding shock.

Total effects, not instrumented buyouts. A funding shock changes several servicing margins at once, including buyout, modification, and liquidation timing, so it is not a valid instrument for any one of them. The primary borrower-level estimand is the total effect of the funding shock on the frozen cohort. Because of displacement within the intermediary, effects are reported across the whole book, weighted both by loan count and by balance. The disclosures follow a loan only while it is in a pool, and Ginnie Mae’s terms of use permit analysis of its MBS, so the outcomes observed here are outcomes of loans in pools; outcomes after repurchase require servicer records.

Model evaluation. Testing Proposition 6 requires outcomes for loans on both sides of the decision. The monthly files record whether loans left in pools later cure, prepay, or reach a claim. With outcomes for repurchased loans from servicer records, or bounds on them, scores trained on earlier cohorts can be evaluated at each issuer’s realized funded share and their ranking compared with the AUC ranking. Where accuracy profiles cross, the model predicts that issuers with different funded shares should prefer different models.

Inference. Standard errors should be clustered at the level at which funding varies. With few clusters, inference should use the wild cluster bootstrap (Cameron, Gelbach, and Miller, 2008), and where parallel trends are invoked, sensitivity should follow Rambachan and Roth (2023) using the full covariance matrix of the event-study coefficients.

7 Conclusion

Debt relief requires information, incentives, and the cash to act. This paper shows that the third shapes the value of the first. An intermediary short of liquidity rations relief by value per dollar, delivers less relief that looks more successful, and values borrower information most when its funded share is intermediate. The value of information peaks where the funded share equals the share of borrowers with good news. It consists of ranking value when liquidity is scarce and becomes screening value as liquidity approaches full funding. The same balance-sheet dependence determines which predictive model a servicer should use, and it makes AUC an average over liquidity states the servicer does not face. Information expands funding only when lenders can verify the technology that produces it. Private selection can move relief away from the borrowers who gain most, and compensation paid after the fact cannot raise the cash that an intermediary whose funding binds commits to relief.

These results suggest a different way of evaluating relief programs and the technology that supports them. The cash that illiquid intermediaries commit to relief responds to liquidity, not to the promised value of relief. The performance of treated loans is a poor guide to the quality of relief when access is rationed. And the benefits of better servicer information depend on the balance sheets of those who hold it. In the Ginnie Mae buyout market from 2019 to 2022, constrained issuers' exercise collapsed in 2020 and, as funding returned, recovered along the ranking they had used before the shock; while cash was scarcest they tilted against the loans that absorbed most of it, and the options they postponed were exercised later on the same ranking. Unconstrained issuers expanded their exercise when value rose, and made most often

the repurchases that constrained issuers left aside, those of loans in forbearance; when value fell in 2022, they cut exercise and began to screen. Testing the remaining predictions, with dated funding variation, outcomes on both sides of the decision, and pre-decision borrower signals, is the natural next step.

Declarations

Declaration of competing interest. From 2019 to 2021 the author was Principal Data Scientist at Rocktop Partners, a mortgage credit investor active in the Ginnie Mae early-buyout market. The firm is not a data source for this paper and had no role in its design, analysis, or writing. The author has no other competing interests to declare.

Data availability. The empirical analysis uses Ginnie Mae’s public loan-level MBS disclosures, together with Freddie Mac’s Primary Mortgage Market Survey (FRED) and state unemployment rates (BLS), both bundled with the code. The loan-level analysis files were assembled from the published disclosures by Patrick Smith and Dongshin Kim for Bandyopadhyay, Kim, and Smith (2026), who shared them, and the issuer charter classification follows Bandyopadhyay (2026d), extended by charter to the largest issuers it does not cover; any researcher can reconstruct the files from the same public source, and the accompanying code rebuilds every reported number from them. Under Ginnie Mae’s terms of use the loan-level files are not redistributed; the package contains the aggregate outputs. The paper uses no proprietary data.

Funding. This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Declaration of generative AI and AI-assisted technologies in the writing process. During the preparation of this work the author used an AI language model (Claude, Anthropic)

to assist with drafting, editing, and code for the numerical examples and the empirical analysis. The author reviewed and edited all output and takes full responsibility for the content of the publication.

References

- Adelino, Manuel, Kristopher Gerardi, and Paul S. Willen, 2013, Why don't lenders renegotiate more home mortgages? Redefaults, self-cures and securitization, *Journal of Monetary Economics* 60, 835–853.
- Agarwal, Sumit, Gene Amromin, Itzhak Ben-David, Souphala Chomsisengphet, Tomasz Piskorski, and Amit Seru, 2017, Policy intervention in debt renegotiation: Evidence from the Home Affordable Modification Program, *Journal of Political Economy* 125, 654–712.
- Aiello, Darren J., 2022, Financially constrained mortgage servicers, *Journal of Financial Economics* 144, 590–610.
- Bandyopadhyay, Arka P., 2026a, Delegated relief: What servicers say, and what statutory eligibility is worth in mortgage forbearance, Working paper.
- Bandyopadhyay, Arka P., 2026b, Inside mortgage loss mitigation: Application completeness, eligibility decisions, and foreclosure risk, Working paper.
- Bandyopadhyay, Arka P., 2026c, What relief buys: Forbearance, modification, and foreclosure outside the CARES Act perimeter, Working paper, SSRN 7427979.
- Bandyopadhyay, Arka P., 2026d, Who can afford to exercise? Non-transferable options and the immobilisation of distressed mortgages, Working paper, SSRN 7427960.
- Bandyopadhyay, Arka P., Dongshin Kim, and Patrick S. Smith, 2026, Adverse selection in mortgage markets: Evidence from Ginnie Mae early buyouts, *Journal of Financial and Quantitative Analysis* 61, 1148–1177.
- Bandyopadhyay, Arka P., and Lilia Maliar, 2026, Reinforcement learning for household finance: Designing policy via responsiveness, *Journal of Economic Dynamics and Control* 182, 105229.
- Berg, Tobias, Valentin Burg, Ana Gombović, and Manju Puri, 2020, On the rise of FinTechs: Credit scoring using digital footprints, *Review of Financial Studies* 33, 2845–2897.
- Blackwell, David, 1953, Equivalent comparisons of experiments, *Annals of Mathematical Statistics* 24, 265–272.
- Boyle, Glenn W., and Graeme A. Guthrie, 2003, Investment, uncertainty, and liquidity, *Journal of Finance* 58, 2143–2166.
- Buchak, Greg, Gregor Matvos, Tomasz Piskorski, and Amit Seru, 2018, Fintech, regulatory arbitrage, and the rise of shadow banks, *Journal of Financial Economics* 130, 453–483.
- Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller, 2008, Bootstrap-based improvements for inference with clustered errors, *Review of Economics and Statistics* 90, 414–427.
- Cherry, Susan F., Erica Xuewei Jiang, Gregor Matvos, Tomasz Piskorski, and Amit Seru, 2021, Government and private household debt relief during COVID-19, *Brookings Papers on Economic Activity* 2021 (Fall), 141–221.
- DeMarzo, Peter, and Darrell Duffie, 1999, A liquidity-based model of security design, *Econometrica* 67, 65–99.
- Froot, Kenneth A., David S. Scharfstein, and Jeremy C. Stein, 1993, Risk management: Coordinating corporate investment and financing policies, *Journal of Finance* 48, 1629–1658.
- Fuster, Andreas, Paul Goldsmith-Pinkham, Tarun Ramadorai, and Ansgar Walther, 2022, Pre-

- dictably unequal? The effects of machine learning on credit markets, *Journal of Finance* 77, 5–47.
- Fuster, Andreas, Matthew Plosser, Philipp Schnabl, and James Vickery, 2019, The role of technology in mortgage lending, *Review of Financial Studies* 32, 1854–1899.
- Ginnie Mae, 2020a, Availability of Pass-Through Assistance Program for participants in Ginnie Mae’s single-family MBS program, All Participants Memorandum 20-03, April 10.
- Ginnie Mae, 2020b, Temporary pooling restrictions on re-performing loans, All Participants Memorandum 20-07, June 29.
- Hand, David J., 2009, Measuring classifier performance: A coherent alternative to the area under the ROC curve, *Machine Learning* 77, 103–123.
- Hernández-Orallo, José, Peter Flach, and Cèsar Ferri, 2012, A unified view of performance metrics: Translating threshold choice into expected classification loss, *Journal of Machine Learning Research* 13, 2813–2869.
- Holmström, Bengt, and Jean Tirole, 1997, Financial intermediation, loanable funds, and the real sector, *Quarterly Journal of Economics* 112, 663–691.
- Holmström, Bengt, and Jean Tirole, 1998, Private and public supply of liquidity, *Journal of Political Economy* 106, 1–40.
- Jiang, Erica Xuwei, Gregor Matvos, Tomasz Piskorski, and Amit Seru, 2020, Banking without deposits: Evidence from shadow bank call reports, NBER Working Paper 26903.
- Kim, You Suk, Steven M. Laufer, Karen Pence, Richard Stanton, and Nancy Wallace, 2018, Liquidity crises in the mortgage market, *Brookings Papers on Economic Activity* 2018 (Spring), 347–428.
- Kim, You Suk, Donghoon Lee, Tess Scharlemann, and James Vickery, 2024, Intermediation frictions in debt relief: Evidence from CARES Act forbearance, *Journal of Financial Economics* 158, 103873.
- Kleinberg, Jon, Himabindu Lakkaraju, Jure Leskovec, Jens Ludwig, and Sendhil Mullainathan, 2018, Human decisions and machine predictions, *Quarterly Journal of Economics* 133, 237–293.
- Kruger, Samuel, 2018, The effect of mortgage securitization on foreclosure and modification, *Journal of Financial Economics* 129, 586–607.
- Liberti, José María, and Mitchell A. Petersen, 2019, Information: Hard and soft, *Review of Corporate Finance Studies* 8, 1–41.
- Maturana, Gonzalo, 2017, When are modifications of securitized loans beneficial to investors? *Review of Financial Studies* 30, 3824–3857.
- Myers, Stewart C., and Nicholas S. Majluf, 1984, Corporate financing and investment decisions when firms have information that investors do not have, *Journal of Financial Economics* 13, 187–221.
- Piskorski, Tomasz, Amit Seru, and Vikrant Vig, 2010, Securitization and distressed loan renegotiation: Evidence from the subprime mortgage crisis, *Journal of Financial Economics* 97, 369–397.
- Rambachan, Ashesh, and Jonathan Roth, 2023, A more credible approach to parallel trends, *Review of Economic Studies* 90, 2555–2591.
- Shaked, Moshe, and J. George Shanthikumar, 2007, *Stochastic Orders* (Springer, New York).
- Stein, Jeremy C., 1997, Internal capital markets and the competition for corporate resources, *Journal*

of Finance 52, 111–133.

Stein, Jeremy C., 2002, Information production and capital allocation: Decentralized versus hierarchical firms, *Journal of Finance* 57, 1891–1921.

A Proofs

Proof of Proposition 1. For any $\lambda \geq 0$ and feasible x ,

$$\int g_i x_i dF = \lambda \int \ell_i x_i dF + \int (g_i - \lambda \ell_i) x_i dF \leq \lambda K + \int [g_i - \lambda \ell_i]_+ dF,$$

which gives weak duality. Let $C(\lambda) = \int \ell_i \mathbf{1}\{g_i > \lambda \ell_i\} dF$ be the cash needed to fund every opportunity whose value per dollar exceeds λ . C is nonincreasing and right-continuous, and it is zero for $\lambda \geq \sup g/\ell$. If $C(0) \leq K$, set $\lambda = 0$ and fund every positive-value opportunity; the bound holds with equality. Otherwise let $\lambda = \inf\{\lambda' : C(\lambda') \leq K\} > 0$. Fund all opportunities with $g_i > \lambda \ell_i$, and fund those with $g_i = \lambda \ell_i$ with the probability that exhausts K ; nonatomicity makes this feasible. Then $\int \ell_i x_i dF = K$, and $x_i = 1$ exactly where $g_i - \lambda \ell_i > 0$, so both inequalities hold with equality and (2) follows. Monotonicity is immediate because a larger budget enlarges the feasible set. For concavity, if x^1 and x^2 are feasible at K_1 and K_2 , their convex combination is feasible at the same combination of budgets and attains the same combination of values. By (2), V is the lower envelope of affine functions of K with slopes λ , so its derivative, where it exists, equals the minimizing λ , and the minimizers are nonincreasing in K . $\square$

Proof of Proposition 2. Under (3), g/ℓ is strictly increasing in q , since its derivative is $(Rk - \omega c)/(k - \omega q)^2 > 0$, so the threshold rule of Proposition 1 is a cutoff in q : $Rq - c > \lambda(k - \omega q)$ if and only if $q > (c + \lambda k)/(R + \lambda \omega)$. Differentiating gives $dq^*/d\lambda$. The cash used at cutoff u is $\int_u^1 (k - \omega q) dF(q)$, with derivative $-(k - \omega u)f(u) < 0$, so a lower K raises the cutoff. The funded share $1 - F(u)$ falls. With $\bar{q}(u) = \mathbb{E}[q \mid q \geq u]$, $\bar{q}'(u) = f(u)[\bar{q}(u) - u]/(1 - F(u)) \geq 0$, strictly when there is mass above u . If q is the success probability under intervention, the law of iterated expectations makes $\bar{q}(u)$ the success rate among funded interventions. $\square$

Proof of Proposition 3. With $\ell_i = k$, the threshold rule funds opportunities in decreasing order of g until either cash runs out or value reaches zero, so the funded share is $\min\{m, \Pr(g >$

0)\} (with randomization at atoms). Parts (a) and (b) follow directly. For (c), Proposition 1 implies that cash deployed equals $\min\{K, \int \ell_i \mathbf{1}\{g_i > 0\} dF\}$, which equals K whenever positive-value cash demand exceeds K . $\square$

Proof of Proposition 4. *Value formula.* Couple $Q = G^{-1}(U)$ with U uniform on $[0, 1]$. For any policy $x(\cdot) \in [0, 1]$ with $\mathbb{E}[x] \leq m$, $\mathbb{E}[(RQ - c)x] \leq \mathbb{E}[(RQ - c)_+ x] \leq \int_{1-m}^1 [RG^{-1}(u) - c]_+ du$, because $[RG^{-1}(u) - c]_+$ is nondecreasing in u and the top- m integral maximizes $\mathbb{E}[h(U)x]$ over such policies for any nonnegative nondecreasing h . The policy $x = \mathbf{1}\{U \geq 1 - m\}\mathbf{1}\{RQ \geq c\}$ attains the bound, which gives (5). Without the signal, all opportunities are identical with value $Rp - c > 0$, so $V_0(m) = m(Rp - c)$.

(i) Writing $[z]_+ = z + [-z]_+$ inside (5) gives $V_G(m) = RT_G(m) - cm + \mathcal{S}(m)$, so $\Delta V = R[T_G(m) - mp] + \mathcal{S}(m)$. Since $\mathbb{E}[Q] = p$, $T_G(1) = p$ and $\mathcal{A}(1) = 0$. $\mathcal{A}'(m) = R[G^{-1}(1 - m) - p]$ is nonincreasing, so $\mathcal{A}$ is concave, and a concave function that vanishes at both endpoints is nonnegative. $\mathcal{S}'(m) = [c - RG^{-1}(1 - m)]_+$ is nonnegative and nondecreasing in m , and it is zero whenever $RG^{-1}(1 - m) \geq c$, which holds for all $m \leq \Pr(RQ \geq c)$.

(ii) $V'_G(m) = [RG^{-1}(1 - m) - c]_+$ is nonincreasing, so V_G is concave; V_0 is linear; hence ΔV is concave. $\Delta V(1) = \mathbb{E}[RQ - c]_+ - (Rp - c) = \mathbb{E}[c - RQ]_+$. For almost every m , $\Delta V'(m) = [RG^{-1}(1 - m) - c]_+ - (Rp - c)$. If $m < m^*$, then $1 - m > G(p)$, so $G^{-1}(1 - m) > p$ and $\Delta V'(m) > 0$. If $m > m^*$, then $1 - m < G(p)$, and because $\Pr(Q = p) = 0$, $G^{-1}(1 - m) < p$. Since $Rp - c > 0$, $[RG^{-1}(1 - m) - c]_+ < Rp - c$, so $\Delta V'(m) < 0$.

(iii) Since $V'(K) = V'_G(m)/k$ with the signal and $V'_0(m)/k$ without it, the cross-difference in marginal values equals $\Delta V'(m)/k$, whose sign is given in (ii).

(iv) For $u > G(p) = 1 - m^*$, $G^{-1}(u) > p$, so $RG^{-1}(u) - c > 0$. Hence $\Delta V(m^*) = R \int_{1-m^*}^1 [G^{-1}(u) - p] du = R \mathbb{E}[(Q - p)\mathbf{1}\{Q > p\}] = R \mathbb{E}[Q - p]_+$, and $\mathcal{S}(m^*) = 0$. Because $\mathbb{E}[Q - p] = 0$, $\mathbb{E}[Q - p]_+ = \mathbb{E}[p - Q]_+ = \frac{1}{2}\mathbb{E}|Q - p|$. $\square$

Proof of Corollary 1. Bayes' rule gives q_H , q_L , and $\pi_H q_H + (1 - \pi_H)q_L = p$. With $Rq_L > c$, $\mathcal{S} \equiv 0$ and $\Delta V = R[T_G(m) - mp]$. For $m \leq \pi_H$, $T_G(m) = q_H m$. For $m \geq \pi_H$,

$T_G(m) = \pi_H q_H + (m - \pi_H) q_L = p - (1 - m) q_L$, so $\Delta V = R(p - q_L)(1 - m)$. Here $m^* = \pi_H$. $\square$

Proof of Proposition 5. With $RQ \geq c$, the intermediary funds the top share m of posteriors, so $x = \mathbf{1}\{U \geq 1 - m\}$ in the coupling $Q = G^{-1}(U)$, and $r = h(G^{-1}(U))$ is nondecreasing in U . (i) $\text{Cov}(x, r) = \mathbb{E}[xr] - \mathbb{E}[x]\mathbb{E}[r] = \int_{1-m}^1 h(G^{-1}(u)) du - m \mathbb{E}[r]$. Its derivative in m is $h(G^{-1}(1-m)) - \mathbb{E}[r]$, which is nonincreasing, so the covariance is concave; it is zero at $m = 0$ and, since the integral equals $\mathbb{E}[r]$ at $m = 1$, at $m = 1$; and its derivative changes sign where $h(G^{-1}(1-m))$ crosses $\mathbb{E}[r]$, that is, at $m = \Pr(r > \mathbb{E}[r])$. For h the identity, the covariance is $T_G(m) - mp = \mathcal{A}(m)/R$, and the crossing is at m^* . (ii) Dividing the covariance by $\Pr(x = 1) = m$ gives $\mathbb{E}[r \mid x = 1] - \mathbb{E}[r]$, and $\mathbb{E}[r \mid x = 1] = m^{-1} \int_0^m h(G^{-1}(1-v)) dv$ is the average of a nonincreasing function over $[0, m]$, hence nonincreasing in m . (iii) If $m \leq m^*$, every funded loan has $Q > p$, so $\Pr(x = 1 \mid Q > p) = m/m^*$ and $\Pr(x = 1 \mid Q < p) = 0$. If $m \geq m^*$, every loan with $Q > p$ is funded and the remaining $m - m^*$ is drawn from the $1 - m^*$ with $Q < p$, so $D = 1 - (m - m^*)/(1 - m^*) = (1 - m)/(1 - m^*)$. $\square$

Proof of Proposition 6. Equation (6) follows from the definitions. (i) Let $t = c/R$ and $h_G(m') = T_G(m') - tm'$. Since $h'_G(m') = G^{-1}(1 - m') - t$ is nonincreasing, h_G is concave and maximized over $[0, m]$ at $\min\{m, \Pr(Q > t)\}$. The top- m integral in (5) keeps exactly the quantiles with $G^{-1}(u) > t$, so $V_G(m) = R \max_{m' \leq m} h_G(m')$. If $T_A \geq T_B$ pointwise, then $h_A \geq h_B$ pointwise and $V_A(m) \geq V_B(m)$ for every m and c . Conversely, with $c = 0$, $V_G(m) = RT_G(m)$, so dominance at every m requires $T_A \geq T_B$. For distributions with equal means, $T_A(m) \geq T_B(m)$ for all m is equivalent to $Q_A \geq_{\text{cx}} Q_B$ (Shaked and Shanthikumar, 2007, Theorem 3.A.5). For a binary state, one experiment is Blackwell more informative than another if and only if its distribution of posteriors is a mean-preserving spread of the other's (Blackwell, 1953). Since $\Gamma_G = T_G/p$, the condition is equivalent to $\Gamma_A \geq \Gamma_B$.

(ii) With $RQ \geq c$ almost surely, $\mathcal{S} \equiv 0$ and $\Delta V_G = \mathcal{A}_G$. Exchanging the order of integration,

$$\int_0^1 \mathcal{A}_G(m) dm = R \int_0^1 [G^{-1}(u) - p] u du = R \text{Cov}(Q, U),$$

where U is the uniform rank in the coupling $Q = G^{-1}(U)$, which randomizes within ties. Let Y be the success indicator, with $\mathbb{E}[Y \mid Q, U] = Q$, so $\text{Cov}(Q, U) = \text{Cov}(Y, U) = p(\mathbb{E}[U \mid Y = 1] - \frac{1}{2})$. Let U' be the rank of an independent draw. Then $\mathbb{E}[U \mid Y = 1] = \Pr(U' < U \mid Y = 1) = p \cdot \frac{1}{2} + (1 - p) \Pr(U_0 < U_1)$, where U_1 and U_0 are the ranks of an independent success and failure. With randomization within ties, $\Pr(U_0 < U_1) = \Pr(Q_0 < Q_1) + \frac{1}{2} \Pr(Q_0 = Q_1) = \text{AUC}_G$. Substituting gives $\text{Cov}(Y, U) = p(1 - p)(\text{AUC}_G - \frac{1}{2})$.

(iii) If $\Gamma_A(m_1) > \Gamma_B(m_1)$ and $\Gamma_A(m_2) < \Gamma_B(m_2)$, then with c small enough that $RQ \geq c$ under both scores, (6) gives strict rankings in opposite directions at m_1 and m_2 , and (ii) gives the AUC ranking of the averages. $\square$

Proof of Proposition 7. (a) Each funded loan requires internal cash $k - Lp > 0$ whatever its posterior, so $m \leq K_0/(k - Lp)$. Since $RQ > c$, the budget binds.

(b) Given a per-loan advance α , the intermediary receives α now and repays L upon success, so a funded loan contributes $(R - L)Q - c + \alpha$ and uses internal cash $k - \alpha$. Because $R > L$, the intermediary funds the highest posteriors and, since each marginal loan has positive value, exhausts the budget: $m(k - \alpha) = K_0$. A competitive lender that knows G and observes the funded set breaks even at $\alpha = L \mathbb{E}[Q \mid \text{top } m] = L T_G(m)/m$. Combining gives $km - LT_G(m) = K_0$. The left side has derivative $k - LG^{-1}(1 - m) \geq k - L > 0$, so the solution is unique. At the solution, the marginal loan has $Q_m \leq \alpha/L$, so $(R - L)Q_m - c + \alpha \geq RQ_m - c > 0$, which confirms that the intermediary funds up to the budget. Under loan-level disclosure, the advance is LQ_i and each loan uses internal cash $k - LQ_i$, which is lowest for the highest posteriors; funding the top share m uses $km - LT_G(m)$, which gives the same m_R .

(c) $T_G(m) \geq mp$, with strict inequality for $m \in (0, 1)$ when Q is nondegenerate (Proposition 4(i)). Hence $km - LT_G(m) \leq m(k - Lp)$, so any m feasible under (a) is feasible under (b), and m_H is strictly slack under (b). Because $K_0 < k - Lp$, both shares are below one.

(d) Financing is fair at the pool level under (b), so private value is $V_G(m_R)$. Under (a), the intermediary receives Lp per loan but repays L with probability equal to its own posterior, a

transfer of $L[T_G(m_H) - m_{HP}] \geq 0$ to the lender. Since V_G is nondecreasing and $m_R \geq m_H$, $V_G(m_R) \geq V_G(m_H) \geq V_G(m_H) - L[T_G(m_H) - m_{HP}]$. $\square$

Proof of Proposition 8. Without the signal, the intermediary cannot distinguish loans within the cell, so any funded share m is independent of Q and borrower surplus is $m \mathbb{E}[\beta(Q)]$. With it, the intermediary funds the top share m , and borrower surplus is $\int_{1-m}^1 \beta(G^{-1}(u)) du$. If β is nonincreasing, $\beta \circ G^{-1}$ is nonincreasing in u , and its average over $[1-m, 1]$ is at most its average over $[0, 1]$, which gives (i); the inequality is strict for $m \in (0, 1)$ unless $\beta \circ G^{-1}$ is almost everywhere constant. The same argument applied to $\bar{w}$ gives (ii), and private value rises by Proposition 4. For (iii), the planner can replicate random allocation, so its optimum with the signal is at least $m \mathbb{E}[\bar{w}(Q)]$. $\square$

Proof of Proposition 9. (a) A payment s raises every g_i by s , which preserves the ranking and weakly enlarges the positive-value set, so the budget still binds and the funded set is the same top share K/k . With a positive density of values at the margin, the multiplier is unique, and the marginal opportunity's value g^* satisfies $g^* + s = \lambda'k$ against $g^* = \lambda k$ before, so $\lambda' = \lambda + s/k$. An injection raises K/k one-for-one while the budget binds.

(b) The budget binds before the payment and the positive-value set weakly expands, so cash deployed remains K . Let x and x' be optimal at payments $s < s'$, with values $\Gamma(x) = \int g_i x_i dF$ and counts $N(x) = \int x_i dF$. Optimality gives $\Gamma(x') + s'N(x') \geq \Gamma(x) + s'N(x)$ and $\Gamma(x) + sN(x) \geq \Gamma(x') + sN(x')$. Adding the two yields $(s' - s)[N(x') - N(x)] \geq 0$, so the count weakly rises. It rises only if the composition of funded actions changes, and the cash constraint then requires that actions with larger cash needs be replaced by actions with smaller ones. A payment $\sigma \ell_i$ changes the ranking ratio from g_i/ℓ_i to $g_i/\ell_i + \sigma$, which preserves the order, so the allocation is unchanged.

(c) With funded share m and marginal posterior $q(m) = G^{-1}(1 - m)$, external borrowing is $B = [mk - K]_+$ and the first-order condition at an interior optimum with $B > 0$ is $Rq(m) - c + s = k\Phi'(mk - K)$. The left side is decreasing in m and the right side is

increasing, so m rises with s and, holding $s = 0$, with K . For a payment s that yields share m_s , the injection $\Delta K = m_s k - K - (\Phi')^{-1}((Rq(m_s) - c)/k)$ yields the same share, provided $(Rq(m_s) - c)/k \geq \Phi'(0)$; otherwise the injection optimum is a corner with $B = 0$ and $\Delta K = m_s k - K$. Under the payment, $B = m_s k - K$ exceeds its initial value because m_s rises. Under the injection, $\Phi'(B) = (Rq(m_s) - c)/k$ is below its initial value because $q(m_s)$ falls, so B falls, to zero in the corner case. $\square$

Figures and Tables




Figure 1. Rationing relief under a funding constraint.

The figure plots the solution of problem (1) with $g(q) = 1.2q - 0.2$, $\ell(q) = 0.3 - 0.1q$, and posteriors uniform on $[0, 1]$. Liquidity is scaled by $\bar{K}$, the cash needed to fund every positive-value opportunity. Panel A shows the share of opportunities funded and the success rate among funded opportunities. Panel B shows the shadow value of cash λ (left axis) and the posterior cutoff q^* (right axis). All quantities are exact model calculations; none is estimated.

Table 1. Rationing as Liquidity Increases

	Cutoff		Shadow value	Funded	Private	Success
$K/\bar{K}$	K	q^*	λ	share	value	among funded
0.100	0.020	0.902	4.204	0.098	0.092	0.951
0.300	0.060	0.718	2.898	0.282	0.234	0.859
0.600	0.121	0.467	1.422	0.533	0.363	0.733
1.000	0.201	0.167	0.000	0.833	0.417	0.583

The table reports the model of Figure 1 at four liquidity levels. Private value and liquidity are per unit mass of opportunities. Success is the success probability among funded interventions. An independent linear program with 4,000 posterior types reproduces private value to within 1.3×10^{-8} at 17 liquidity levels.




Figure 2. The value of information is hump-shaped in liquidity.

Posteriors are distributed $\text{Beta}(1.5, 1)$, with mean $p = 0.6$; $R = 1$ and $c = 0.35$, so 20.7% of borrowers are unprofitable on the finer information. Panel A plots the allocation value of the signal, $\Delta V(m) = V_G(m) - V_0(m)$, and its ranking and screening components (Proposition 4). The peak is at $m^* = \Pr(Q > p) = 0.535$ with value $\mathbb{E}[Q - p]_+ = 0.112$. Panel B plots the marginal value of liquidity per unit of fundable share with and without the signal; the curves cross at m^* .





Figure 3. Which score is worth more depends on liquidity.

Two calibrated scores with mean $p = 0.5$, $R = 1$, and $c = 0.1$. Score A assigns posterior 1 to 20% of borrowers and 0.375 to the rest; score B assigns 0.8 and 0.2 to halves. Panel A plots each score's allocation value by fundable share. Vertical lines mark 7.2% and 86.2%, the nonbank and depository buyout exercise rates in June 2020 in the sample of Section 6; the signal structures are hypothetical. Panel B plots the cumulative accuracy profiles; the allocation value in Panel A is Rp times the vertical distance between each profile and the diagonal.

Table 2. Value of Two Scores at Different Fundable Shares

Fundable share m	Score A (AUC = 0.70)	Score B (AUC = 0.80)	Ratio A/B
0.072	0.0360	0.0216	1.67
0.250	0.0938	0.0750	1.25
0.500	0.0625	0.1500	0.42
0.750	0.0313	0.0750	0.42
0.862	0.0173	0.0414	0.42

Allocation value per loan in the model of Proposition 6, with the scores of Figure 3. The first and last rows set the fundable share to the nonbank and depository buyout exercise rates in June 2020 in the sample of Section 6. Averaged over fundable shares drawn uniformly from $[0, 1]$, the values are 0.05 and 0.075, equal to $Rp(1 - p)(\text{AUC} - \frac{1}{2})$ as in equation (7).

**Figure 4.** Rationing along a fixed ranking: Ginnie Mae buyout decisions, 2019–2022.

Panel A plots the share of vested buyout options exercised within three months of vesting, by vesting month, at depository and nonbank issuers. Panel B plots, for nonbanks, the relative response of exercise to the ranking index: the within-cell covariance of exercise with the index, divided by its variance and by the funded share, where a cell is an issuer, vesting month, and forbearance status. The fixed-ranking line holds the latent weight on the index at its March 2019–February 2020 estimate from equation (11) and matches each cell’s funded share; it uses no information on which loans were exercised outside that window. The random-rationing benchmark has each issuer fund, along the same ranking, the options it would have funded at its own pre-period share and exercise a random fraction of them that matches its observed funded share. Shading marks March–June 2020 (freeze) and July–December 2020 (release); the dotted line marks the start of 2022.

Table 3. Rationing along a Fixed Ranking, 2019–2022


	Pre-shock	Freeze	Release	Recovery		Rate rise
	3/19–2/20	3/20–6/20	7/20–12/20	1/21–6/21	7/21–12/21	1/22–9/22
<i>Panel A. Nonbank issuers</i>						
Options	175,716	283,048	357,497	142,549	118,289	193,853
Funded share (%)	40.0	8.3	15.9	32.4	37.1	7.2
Latent weight on index ^a	1.45 (0.29)	−0.49 (0.31)	−0.09 (0.24)	−0.07 (0.21)	−0.19 (0.13)	0.60 (0.20)
Level response	0.201	0.043	0.107	0.140	0.136	0.072
Relative response	0.50	0.52	0.67	0.43	0.37	0.99
Fixed ranking	0.50	0.75	0.71	0.45	0.40	0.80
Random rationing	0.46	0.40	0.47	0.37	0.35	0.48
Actual minus fixed	–	−0.23 (0.25)	−0.03 (0.11)	−0.02 (0.05)	−0.04 (0.04)	0.19 (0.11)
Actual minus random	–	0.11 (0.16)	0.20 (0.08)	0.06 (0.05)	0.02 (0.03)	0.52 (0.11)
Same, change from pre	–	0.07 (0.15)	0.16 (0.08)	0.02 (0.05)	−0.02 (0.03)	0.48 (0.11)
<i>Panel B. Depository issuers</i>						
Options	95,934	189,061	183,529	52,483	35,943	47,870
Funded share (%)	80.5	85.3	79.2	79.4	81.0	51.9
Latent weight on index ^a	0.13 (0.03)	−0.35 (0.16)	−0.12 (0.14)	0.11 (0.17)	−0.06 (0.12)	3.89 (1.86)
Level response	0.010	−0.007	0.000	0.013	0.004	0.177
Relative response	0.01	−0.01	0.00	0.02	0.00	0.34
Fixed ranking	0.01	0.00	0.01	0.01	0.01	0.03
Random rationing	0.01	0.00	0.00	0.01	0.01	0.01
Actual minus fixed	–	−0.01 (0.02)	−0.01 (0.02)	0.01 (0.02)	0.00 (0.01)	0.31 (0.12)
Actual minus random	–	−0.01 (0.02)	0.00 (0.02)	0.01 (0.02)	0.00 (0.01)	0.33 (0.13)
Same, change from pre	–	−0.01 (0.02)	0.00 (0.02)	0.01 (0.02)	0.00 (0.01)	0.33 (0.13)

The sample is 1,922,668 Ginnie Mae buyout options that vested from January 2019 to September 2022 at 120 depository and nonbank issuers; exercise is repurchase within three months of vesting. The six periods begin in March 2019, so options that vested in January and February 2019 enter the sample and the index diagnostics but no column. The ranking index is the linear predictor of a logit of exercise on hard loan characteristics with issuer-month fixed effects, estimated on nonbank decisions in 2016–2018 and standardized over the sample. Cells are issuer by vesting month by forbearance status. ^aThe pre-shock column reports the level; later columns report the change from it. The latent weight is the coefficient on the index in a logit with cell fixed effects and period interactions; standard errors are clustered by issuer. The level response is the within-cell covariance of exercise with the index divided by its variance; the relative response divides it by the funded share. Fixed-ranking predictions hold the latent weight at its pre-shock estimate and match each cell’s funded share; in the pre-shock period they equal the actual responses by construction, because the logit’s score equations match the covariance of exercise with the index. The random-rationing benchmark has each issuer fund, along the same ranking, the options it would have funded at its own pre-shock share and exercise a random fraction of them that matches its observed funded share; it lies below the actual response in the pre-shock period because months in which an issuer’s share falls below its own average are scaled down, and the last row therefore compares changes from the pre-shock period. Standard errors for differences are from 200 bootstrap draws over issuers that re-estimate the pre-shock weight.

Table 4. The Price of Cash in the Ranking: The Weight on Loan Balance

	Pre-shock	Freeze	Release	Recovery		Rate rise
	3/19–2/20	3/20–6/20	7/20–12/20	1/21–6/21	7/21–12/21	1/22–9/22
<i>Nonbanks</i>						
Weight on log balance	0.19 (0.11)	−0.32 (0.06)	0.07 (0.06)	0.08 (0.11)	−0.10 (0.14)	−0.06 (0.13)
Weight on ranking index	1.52 (0.30)	−0.62 (0.33)	−0.03 (0.25)	0.00 (0.26)	−0.21 (0.17)	0.60 (0.19)
<i>Depositories</i>						
Weight on log balance	−0.01 (0.01)	−0.02 (0.03)	−0.05 (0.06)	−0.06 (0.06)	−0.06 (0.04)	−0.15 (0.16)
Weight on ranking index	0.12 (0.03)	−0.36 (0.16)	−0.15 (0.11)	0.08 (0.15)	−0.09 (0.11)	3.95 (2.33)

Coefficients from a logit of exercise within three months of vesting on the standardized ranking index and standardized log unpaid balance, with cell fixed effects (issuer by vesting month by forbearance status) and interactions of both regressors with period indicators. The pre-shock column reports levels; later columns report changes from the pre-shock level. Splitting the freeze into March–April, when the forbearance flag is not yet observed, and May–June, when it enters the cell definition, gives changes in the weight on balance of -0.31 (0.15) and -0.32 (0.08) for nonbanks and -0.04 (0.05) and -0.02 (0.04) for depositories. Standard errors are clustered by issuer. The sample is as in Table 3.

**Figure 5.** Postponed exercise and the switch to screening.

Panel A plots the share of options exercised by each number of months after vesting, observed through December 2022, for options that vested at nonbanks and depositories in March 2019–February 2020 (dashed) and March–June 2020 (solid). Panel B plots the level response of exercise within three months to the ranking index, the within-cell covariance of exercise with the index divided by its variance, by vesting month. Shading marks March–June 2020 (freeze) and July–December 2020 (release); the dotted line marks the start of 2022.

Table 5. Postponed Options Were Exercised Later on the Same Ranking

Vesting period	Nonbanks			Depositories		
	Options	Later (%)	Weight	Options	Later (%)	Weight
3/2019–2/2020	92,921	5.2	0.56 (0.15)	13,334	23.6	0.11 (0.03)
3/2020–6/2020	240,382	26.4	1.05 (0.25)	21,004	24.5	0.39 (0.24)
7/2020–12/2020	275,410	27.0	0.80 (0.20)	32,859	23.4	0.21 (0.16)
1/2021–6/2021	80,916	20.2	0.46 (0.14)	8,272	17.6	−0.07 (0.16)

The sample is options not exercised within three months of vesting and still in their pools at the end of the third month. Later is the share exercised in the fourth through twelfth months after vesting. Weight is the coefficient on the ranking index in a logit of later exercise with cell fixed effects (issuer by vesting month by forbearance status); standard errors clustered by issuer are in parentheses. For comparison, the pre-shock weight on exercise within three months is 1.45 for nonbanks (Table 3).

Table 6. Repurchases Depositories Made and Nonbanks Did Not: Forbearance, May–June 2020

	Nonbanks	Depositories
Share of options in forbearance (%)	90	88
Exercised within three months, in forbearance (%)	3.9	87.6
Exercised within three months, not in forbearance (%)	38.6	72.1
Difference within issuer-month (pp)	−27.2 (10.3)	13.2 (3.0)
Same, within index decile (pp)	−26.6 (10.0)	13.2 (3.0)
Exercised by June 30, 2020, in forbearance (%)	1.1	76.4
Exercised by June 30, 2020, not in forbearance (%)	37.1	66.8
Exercised by December 2022, in forbearance (%)	38.5	91.0
Issuers	77	36

Buyout options that vested in May and June 2020, the first months in which the disclosures flag forbearance. The within-issuer-month difference is the coefficient on a forbearance indicator in a linear probability model of exercise with issuer-month fixed effects; the next row adds indicators for deciles of the ranking index. Standard errors are clustered by issuer.

Table 7. Predictions and Evidence

Prediction	Observable implication	Current evidence
1. Profitable actions go unfunded, ranked by value per dollar (Prop. 1)	Constrained intermediaries forgo or postpone actions whose value has not fallen, and tilt against cash-intensive ones when cash is scarce	Nonbanks' weight on balance fell by 0.32 in the freeze and recovered; depositories' unchanged (Table 4); by June 30, 2020, depositories had repurchased 76.4% of May–June options in forbearance, nonbanks 1.1% (Table 6)
2. Tighter liquidity improves measured performance (Prop. 2)	Cure rate among funded actions rises as the funded share falls, holding the pre-event distribution of borrower states fixed	Selection on the index did not fall as nonbank funded shares fell almost fivefold, and rose, in issuer-level point estimates, at four of the five largest nonbanks and at the small issuers as a group (Table 3; Internet Appendix Table IA.3); cure rates require outcomes after repurchase
3. Liquidity moves constrained intermediaries; value moves unconstrained ones (Prop. 3)	Opposite responses when a common cash drain coincides with a common rise in value; asymmetric responses to a dated value shock	2020: nonbank funded share fell almost fivefold, depository share rose; 2022 value shock: depository share fell from 81.0% to 51.9% and the ranking weight rose at both types (Table 3)
4. Level reliance on information is hump-shaped; relative reliance rises as liquidity contracts (Prop. 5)	Level slope of relief in a ranking variable falls as the funded share falls below the peak; slope relative to the base rate rises	Nonbank relative response was 0.52 in the freeze, between random rationing (0.40) and the fixed ranking (0.75); 0.67 in the release and 0.43 and 0.37 in 2021 against fixed-ranking values of 0.71, 0.45, and 0.40; postponed options exercised on the same ranking (Table 5)
5. Information and liquidity switch from complements to substitutes at m^* (Props. 4 and 5)	Funding shock $\times$ prior signal interaction positive below and negative above $\hat{m}^*$	Screening at full funding: depositories' level response rose from 0.010 to 0.177 when the 2022 value shock created negative-value options (Figure 5); funding interaction untested, equation (12)
6. Model rankings depend on liquidity (Prop. 6)	Out-of-sample value of scores at issuer-specific funded shares departs from the AUC ranking where accuracy profiles cross	Untested; requires outcomes on both sides of the decision
7. Recognized information expands funding (Prop. 7)	Facility advance rates respond to verifiable model adoption or data sharing, not to internal adoption alone	Untested
8. Private selection can lower borrower surplus (Prop. 8)	Borrower gains from relief decline in the intermediary's predicted recovery	Self-cure evidence (Adelino, Gerardi, and Willen, 2013); direct test requires borrower outcomes
9. Ex post compensation is inframarginal under binding funding (Prop. 9)	Per-action incentives leave the cash committed to relief unchanged at constrained servicers, shift relief toward cheaper actions, and raise facility draws where funding is elastic	Staff capacity, the analogous constraint in Section 5, limited HAMP modifications (Agarwal et al., 2017); untested

Internet Appendix for

“The Financing of Debt Relief: Intermediary Liquidity and the Value of Borrower Information”

IA.1 Auxiliary Results

IA.1.1 Blackwell improvements under a fixed funding technology

If one experiment is a garbling of another and the intermediary can ignore the finer signal, every policy feasible under the coarser experiment remains feasible under the finer one. With the funding technology held fixed, private value before information costs cannot fall. The statement concerns the intermediary’s own information. A mandatory disclosure policy can change outside financing terms and equilibrium beliefs, and Proposition 7 shows that what lenders know about the technology, not the technology alone, determines borrowing capacity. The claim that better information always expands funding conflates these cases.

IA.1.2 Why a flatter decision slope is not a unique liquidity signature

The action rule in Proposition 1 is a threshold, not a linear response, so the slope of relief in a ranking variable depends on where the threshold sits. Three cases illustrate. First, a binding funding constraint preserves the ranking and funds less of it: the level slope falls toward zero as liquidity contracts below the peak, and the slope relative to the base rate rises (Proposition 5). Second, a value shock that shifts every opportunity’s value by the same amount also preserves the ranking. An unconstrained intermediary then funds the top of the same ranking, and at the same funded share it produces the same level and relative slopes as rationing. Third, a value shock that compresses values across the ranking variable, such as a restriction on redelivery that bears most heavily on the loans with the largest redelivery

gains, can flatten both slopes without any binding cash constraint. Slopes alone therefore separate funding from value only against the third alternative. Against the second, what separates them is the response to dated shocks of each kind (Proposition 3): the funded share of a constrained intermediary moves with liquidity and not with value, and that of an unconstrained one does the reverse.

IA.1.3 Staff capacity and funding

In (10), if every feasible action satisfies $d_{ia} = \varrho \ell_{ia}$ for a common $\varrho > 0$, then $g_{ia} - \lambda \ell_{ia} - \mu d_{ia} = g_{ia} - (\lambda + \varrho \mu) \ell_{ia}$, and choices depend only on the combined multiplier. Action data alone cannot then separate funding scarcity from staffing scarcity. Separation requires actions whose cash and processing requirements vary independently, such as a buyout, which is cash-intensive but processing-light, and a documented modification, which is processing-intensive and requires relatively little cash, or a direct measure of one of the two constraints.

IA.1.4 Heterogeneous cash needs and the value of information

Proposition 4 fixes the cash need. With heterogeneous needs, the intermediary ranks by $(RQ - c)/\ell$, and the value of a signal depends on its joint distribution with cash needs. Two observations are immediate. If cash needs are independent of the signal, the intermediary still ranks by Q within each cash-need class, so the funded set in each class is a top share of that class's posteriors. The signal changes the common shadow value of cash, however, and with it the funded share of each class, so the aggregate value of the signal need not be single-peaked in K ; Proposition 4 applies to each class holding its funded share fixed. If cash needs fall with the posterior, as in specification (3) with $\omega > 0$, the signal also identifies cheaper actions, and part of its value is a saving of cash. That component operates only through the budget, so it vanishes when the budget is slack.

IA.2 Dynamic Extension

Let $\mathbf{z}_t$ collect each loan’s posterior, legal status, eligibility, and contractual cash-flow state. A finite-horizon intermediary chooses a vector of actions subject to current cash and staff constraints:

$$V_t(K_t, \mathbf{z}_t) = \max_{\mathbf{a}_t \in \mathcal{F}(K_t, \mathbf{z}_t)} \mathbb{E}_t[\Pi_t(\mathbf{a}_t, \mathbf{z}_t) + \delta V_{t+1}(K_{t+1}, \mathbf{z}_{t+1})]. \quad (\text{IA1})$$

Cash evolves through collected payments, advances, repurchases, operating costs, and financing. Information acquisition uses operating resources and shifts future posteriors. Transitions depend on the action and describe cure, renewed delinquency, completed modification, and liquidation. The static model is a one-period version that uses continuation-adjusted values.

Three features of the dynamic problem reinforce the static results. First, cash spent today reduces K_{t+1} one-for-one while its return arrives with delay, so the shadow value of cash tends to be highest early in a delinquency wave, when advance obligations accumulate. Second, waiting has option value when information arrives over time, so a constrained intermediary can rationally postpone action on loans whose posteriors are near the threshold. Liquidity can therefore affect the timing of relief more than the relief eventually granted, which is why the empirical design measures outcomes at fixed horizons. Third, if the budget binds in each period and an action ties up its cash for τ periods before it is returned, the action uses the budget in each of those periods, so its cost in scarce cash is $\sum_{s < \tau} \delta^s \lambda_{t+s} \ell$, approximately $\tau \lambda \ell$ when multipliers are stable and discounting is small. Ranking by value per dollar then becomes ranking by value per dollar-period: a constrained intermediary passes over actions that tie up cash longest, even when they are worth more. Bandyopadhyay and Maliar (2026) design dynamic, borrower-responsive household-finance policy by reinforcement learning. The extension here adds a budget shared across the portfolio, which creates interference across loans, so that the simulated value of a learned policy is not a causal estimate of its effect if a servicer adopted it institution-wide.

IA.3 Data Construction Protocol

IA.3.1 Sources and the definition of action

Ginnie Mae publishes monthly loan-level disclosure files for single-family pools. In these files, a loan removed from a pool under removal reason 2 (repurchase of a delinquent loan) identifies an exercised buyout. The files used in Section 6 were assembled from the disclosures for Bandyopadhyay, Kim, and Smith (2026): a vesting-date panel with twelve months of outcomes and the final status through December 2022, and the monthly histories of every loan that reached three missed payments. Tests of Propositions 4–6 require three additions: dated measures of issuer funding, a pre-decision ranking variable, and outcomes at fixed horizons after the decision.

IA.3.2 Canonical fields

Table IA.1. Canonical Fields of the Data-Construction Protocol

Field	Definition and verification requirement
<code>loan_id</code>	Stable identifier; document whether it survives repurchase, servicing transfer, and redelivery.
<code>report_date</code>	Economic as-of date, distinct from the file’s publication date.
<code>issuer_id</code>	Issuer of record at the decision date, with a dated crosswalk for mergers and transfers.
<code>pool_id</code>	Security at the date; a loan can pass through more than one pool.
<code>eligible_date</code>	First date the action became available under the rule then in force, not the date the action was taken.
<code>buyout_date</code>	Validated by removal reason; disappearance from a pool alone is not a buyout.

Field	Definition and verification requirement
<code>upb, coupon</code>	Unpaid balance and note rate at the decision date.
<code>funding_capacity</code>	Usable cash and undrawn committed facility capacity, measured before the decision; gross facility size is not availability.
<code>cash_requirement</code>	Current cash required by the action net of immediately draw-able financing; later proceeds are recorded separately.
<code>signal_time</code>	Time stamp showing that the ranking information preceded the decision.
<code>action_stage</code>	Contact, application, determination, offer, acceptance, execution, and completion as separately dated stages.
<code>outcome_date</code>	Date of a verified outcome, distinguishing absence of evidence from evidence of no event.

IA.3.3 Population and attrition

Report, in order, all source loans; loans in the date range; loans observed before eligibility; eligible loans; loans with a classified issuer; loans with funding data; loans with a usable pre-decision signal; and loans with observed outcomes at each horizon. Give both loan and issuer counts at each step, and analyze missing funding or signal data as selection into the sample. Revisions to the same loan-month are not independent observations, and a servicing transfer is a transition rather than a new delinquency spell. End-of-sample censoring and loss of coverage are different events, and a loan that leaves a Ginnie Mae pool can remain outstanding and delinquent outside it.

IA.3.4 Estimands and outcome hierarchy

The primary institution-level estimand is the total effect of a funding change on delivered resolution within an h -month horizon for a cohort frozen before the change, allowing actions

on other loans to respond. The allocation estimand is how the funding change shifts resolution across strata of pre-event information and cash requirement. The principal outcome is a completed, documented intervention within a common horizon; offers, reviews, processing time, repurchase, and redelivery are secondary and mechanism outcomes. Pending foreclosure, completed foreclosure, bankruptcy, and redelivery are distinct states and are not interchangeable welfare measures. Horizons are fixed before outcomes are examined, and confirmatory tests are distinguished from exploratory heterogeneity.

IA.4 Empirical Details

IA.4.1 Sample, outcome, and issuer types

Each observation is a buyout option in the month the loan first reaches three missed payments, from January 2019 to September 2022, in single-family level-payment and reperforming pools of both Ginnie Mae programs. The construction drops loans outside the 50 states and the District of Columbia, loans first paying before 2003, loans with a single delinquent month surrounded by current months, loans with missing or implausible balance or loan-to-value ratio, and loans that cannot be matched to state house prices or Treasury rates. Exercise is removal under reason 2 before any cure, within three months of vesting; the file also records the month of removal within twelve months and the final removal reason through December 2022. Issuer charter types follow Bandyopadhyay (2026d), which classifies the issuers in Ginnie Mae’s 2020 issuer file and the largest unclassified issuers of 2019–2020 by charter. I extend the hand classification by charter to the forty largest unclassified issuers of 2019–2022; unclassified issuers, housing finance agencies, and technology-based lenders are excluded from the main sample.

IA.4.2 Ranking index

The index is estimated on the 298,723 nonbank options that vested from 2016 to 2018 in the 649 issuer-months in which some but not all options were exercised, by a logistic regression of exercise on standardized characteristics and a full set of issuer-month indicators (inverse regularization strength 10). The characteristics are the note rate minus the monthly mean of Freddie Mac’s 30-year mortgage rate in the vesting month, the change in the ten-year Treasury rate since origination, log loan age, log balance, credit score and a missing indicator, current loan-to-value ratio, judicial-foreclosure state, VA and rural-housing indicators, pool issue-type indicators, first-time-buyer status, refinance and modification indicators, prior missed payments in the preceding year, and the annual insurance premium. Standardization and imputation use the training sample. The within-issuer-month AUC, averaged over cells with at least 50 options and both outcomes and weighted by cell size, is 0.73, 0.68, 0.67, and 0.78 for nonbank decisions in 2019, 2020, 2021, and 2022, and 0.54, 0.50, 0.53, and 0.77 for depository decisions. The mortgage spread alone has within-cell AUCs of 0.70, 0.65, 0.64, and 0.79 for nonbanks. Table IA.2 uses four alternatives: the same logit with the spread over the ten-year Treasury rate entered piecewise linearly with kinks at 2, 3, and 4 points, which ranks 2020 decisions less well because almost every 2020 option lies above the first kink; a gradient-boosted ranker (LightGBM, at most 15 leaves, at least 1,000 observations per leaf, learning rate 0.03, early stopping on held-out issuer-months) fit to the same decisions with each cell’s log-odds of exercise as a fixed offset; the conditional logit fit to depository decisions; and a gradient-boosted classifier fit to pooled depository decisions without cell effects, the construction of an earlier version of this paper, which ranks nonbank decisions within cells poorly (within-cell AUC 0.54 in 2019), because depositories exercised almost every option and their decisions reveal little about the order in which a constrained issuer funds options.

IA.4.3 Estimation

The logit with cell fixed effects is estimated by block coordinate ascent: given the slope, each cell's intercept solves its score equation; given the intercepts, the slope takes a Newton step on the concentrated likelihood. Cells in which every option or no option was exercised carry no information about the slope and are dropped from that estimation. Standard errors use the cluster-robust sandwich on the concentrated score, clustered by issuer. For the fixed-ranking predictions, the intercept of each cell is chosen by bisection so that the mean predicted probability equals the observed funded share, holding the slope at its pre-shock estimate; cells with a funded share of zero or one have a predicted response of zero, as they do in the data. For the random-rationing benchmark, the fixed-ranking probabilities are computed at the issuer's own pre-period funded share and the cell's covariance is scaled by the ratio of the observed funded share to that share (or set equal to the fixed-ranking value when the observed share is higher). Level and relative responses aggregate within-cell covariances and variances across the cells of a period. Bootstrap standard errors resample issuers with replacement, re-estimate the pre-shock slope in each of 200 draws, and recompute all predictions. Cells are split by forbearance status because the index cannot rank options in and out of forbearance against each other: it is estimated before forbearance existed, and within issuer-months nonbanks exercised options on loans in forbearance 27.2 percentage points less often than other options (Table 6). Before May 2020 the flag does not exist and the cells are issuer-months. The bootstrap resamples the 83 nonbank issuers and the 37 depositories with replacement. In the files, a removal is recorded in the disclosure month in which the loan leaves the pool, and the vesting-month record is month zero, so exercise within three months covers months zero through three; exercise by June 30, 2020 is a removal recorded in the May or June 2020 file. Loans recorded as reperforming or still delinquent at the end of the twelve-month window (codes 8 and 9 in the construction) remain in the queue risk set.

IA.4.4 Robustness

Table IA.2 repeats the freeze comparison for nonbanks with different ranking variables, decision cells, exercise windows, and samples. All indices are standardized over the depository and nonbank options of January 2019 to September 2022, as in Table 3. With issuer-month cells, the prediction treats options in and out of forbearance as ranked by the same index, although nonbanks funded almost none of the former. In every row that ranks on nonbank decisions, the freeze response lies at or above the random-rationing benchmark and at or below the fixed-ranking prediction, within a few hundredths, and no row overturns the conclusion that the freeze months alone do not separate them. The index estimated from depository decisions, which ranks nonbank decisions poorly, puts the response below both. The standard errors in this table come from 100 draws with a different seed from the 200 draws behind Table 3, which is why the baseline row's standard errors differ slightly.

Table IA.2. Robustness of the Freeze Test, Nonbank Issuers

Specification	Latent weight		Funded	Relative response		Freeze: actual minus	
	Pre	Change, freeze		Pre	Freeze	Fixed	Random
Baseline	1.45	−0.49 (0.31)	8.3	0.50	0.52	−0.23 (0.26)	0.11 (0.13)
Issuer-month cells	1.45	−0.50 (0.33)	8.3	0.50	0.63	−0.28 (0.28)	0.16 (0.17)
Treasury-spread index, piecewise	2.48	0.03 (1.08)	8.3	0.74	1.14	0.01 (0.47)	0.34 (0.42)
Boosted ranking index	0.99	0.01 (0.23)	8.3	0.34	0.47	0.00 (0.14)	0.18 (0.08)
Index from depository decisions	0.81	−0.49 (0.12)	8.3	0.30	0.17	−0.24 (0.12)	−0.07 (0.06)
Depository exercise classifier	0.07	0.07 (0.13)	8.3	0.03	0.07	0.03 (0.08)	0.05 (0.07)
Mortgage spread only	1.20	−0.41 (0.39)	8.3	0.44	0.50	−0.25 (0.43)	0.12 (0.22)
Exercise within 12 months	1.42	−0.35 (0.23)	30.7	0.47	0.49	−0.13 (0.11)	0.03 (0.05)
Excluding the largest nonbank	1.28	−0.22 (0.29)	10.6	0.48	0.51	−0.09 (0.16)	0.16 (0.17)
Nonbanks and technology-based lenders	1.41	−0.45 (0.29)	8.3	0.51	0.53	−0.21 (0.17)	0.11 (0.14)
Loans not in forbearance	1.45	−0.34 (0.36)	25.1	0.50	0.34	−0.08 (0.17)	0.01 (0.10)

Pre is March 2019–February 2020 and freeze is March–June 2020; the funded share is for the freeze, in percent. The last two columns report the freeze relative response minus the fixed-ranking prediction and minus the random-rationing benchmark, constructed as in Table 3, with standard errors from 100 bootstrap draws over issuers. Standard errors for the latent weight are clustered by issuer.

IA.4.5 Issuer-level results

Table IA.3 reports the freeze and recovery comparisons for the five nonbanks with the most options in the pre-shock period, among issuers that exercised at least 5% of them, and for the remaining nonbanks as a group. Issuers are ranked ex ante, by pre-shock option counts; the 5% floor excludes one large issuer that exercised almost none of its pre-shock options and so has no estimable pre-shock weight (it is in the last row). Each row re-estimates the pre-shock weight on that issuer’s own decisions. The aggregate freeze shortfall from the fixed-ranking

prediction in Table 3 comes from issuer B; the other rows are at or above the prediction, except issuer A, which lies between the two benchmarks. Issuer-level responses carry no standard errors, and those of issuers whose funded share fell below 5% rest on few exercises.

Table IA.3. The Largest Nonbanks in the Freeze and the Recovery

Issuer	Pre-shock, March 2019–February 2020				Freeze, March–June 2020				Recovery, January–June 2021			
	Options	Latent weight	Share	Actual	Share	Actual	Fixed	Random	Share	Actual	Fixed	Random
A	36,265	2.19	53.4	0.69	4.2	1.53	2.06	0.74	54.6	0.70	0.63	0.58
B	33,343	2.04	61.3	0.58	2.4	0.38	2.11	0.61	44.1	0.65	0.78	0.57
C	16,256	0.87	63.1	0.29	25.4	0.50	0.38	0.27	17.0	0.58	0.61	0.31
D	12,054	0.42	71.4	0.10	48.9	0.14	0.14	0.07	88.6	0.04	0.05	0.05
E	5,747	0.52	59.0	0.20	25.7	0.62	0.25	0.17	28.5	0.45	0.26	0.21
Other 78 nonbanks	72,051	0.95	11.3	0.50	0.7	1.10	0.90	0.71	15.9	0.12	0.25	0.24

Nonbank issuers ranked by the number of options that vested in March 2019–February 2020, among issuers that exercised at least 5% of them. Shares are funded shares in percent. Actual is the relative response; Fixed and Random are the fixed-ranking prediction and the random-rationing benchmark, constructed as in Table 3 with the latent weight re-estimated on the issuer’s own pre-shock decisions. The last row pools the other 78 nonbanks.

IA.4.6 Funding pressure within issuer types

Table IA.4 relates exercise to four measures of an issuer’s own cash drain, each computed for the issuer’s whole portfolio from the monthly disclosures: advances outstanding as a share of unpaid balance (in logs and in levels), advances outstanding in months of scheduled principal and interest, and the share of loans in forbearance. Advances are scheduled payments times months delinquent, summed over the issuer’s loans. Each regression includes issuer and month fixed effects, weights issuer-months by the number of options or eligible loans, and clusters standard errors by issuer; the funded share and the relative response are computed from the decision cells of Section 6.2, and the monthly exercise rate from the eligible stock of Table IA.5. In a cell-fixed-effects logit with index-by-month interactions, so that the advance term is identified from cross-issuer differences within a month, the latent weight on the index falls with a nonbank’s own log advance burden (-0.97 per log point, s.e. 0.47) and not significantly with a depository’s (-0.10 , 0.17). Adding log balance and its interaction with the

burden to that logit gives index interactions of -1.11 (0.51) for nonbanks and -0.65 (0.39) for depositories, and the weight on balance does not vary with the burden at either type (-0.06 , s.e. 0.12, and -0.14 , 0.15). Two cross-sectional exposures to the 2020 shock, each issuer’s May 2020 forbearance share and its 2019 state mix weighted by the February–April 2020 rise in state unemployment rates, have small dispersion across issuers (standard deviations of 4.3 and 1.2 percentage points among the 58 nonbanks with at least 100 options in both periods) and are not significantly related to the change in an issuer’s funded share from the pre-shock period to the freeze, although the estimates are imprecise: -3.0 points per point of forbearance share (s.e. 2.3), or -12.8 points per standard deviation of the exposure (9.9), and 16.1 per point of unemployment exposure (10.7), or 19.9 per standard deviation (13.2), for nonbanks, against an aggregate fall of 32 points; the depository coefficients are -1.8 (2.0) and -1.2 (8.4).

Table IA.4. Exercise and the Issuer’s Own Cash Drain

Measure	Nonbanks			Depositories		
	Funded share	Stock rate	Rel. resp.	Funded share	Stock rate	Rel. resp.
Log advances outstanding (% of balance)	-9.2 (3.1)	-4.1 (1.6)	-0.01 (0.29)	-3.2 (3.5)	-1.5 (3.8)	0.09 (0.06)
Advances outstanding (% of balance)	-13.6 (13.4)	-6.1 (5.0)	0.74 (1.02)	-13.2 (12.3)	-4.7 (11.4)	0.12 (0.22)
Advances outstanding (months of P&I)	-6.8 (5.3)	-3.2 (2.2)	0.40 (0.53)	-7.0 (6.3)	-1.2 (5.6)	0.08 (0.11)
Share of portfolio in forbearance (%)	-1.5 (0.8)	-0.7 (0.2)	-0.03 (0.05)	-1.7 (1.2)	-2.1 (0.9)	0.01 (0.01)

Issuer-month regressions, January 2019–September 2022, with issuer and month fixed effects, on issuer-months with at least 20 vesting options (2,358 nonbank and 1,229 depository issuer-months). Funded share is the share of options vesting in the month exercised within three months (in percentage points), weighted by options; stock rate is the monthly repurchase rate on the eligible stock (percentage points), on issuer-months with at least 50 eligible loans and weighted by eligible loans; rel. resp. is the relative response of exercise to the ranking index, defined only for issuer-months with some exercise. Coefficients are per unit of the measure in the row. Standard errors clustered by issuer in parentheses.

IA.4.7 The monthly eligible stock

The vesting-level test treats the decision at vesting. An issuer can exercise the option in any month in which the loan remains at least three payments delinquent and in its pool.

Table IA.5 repeats the test on the 12,456,455 loan-months of that stock from January 2019 to September 2022, with the monthly exercise rate as the funded share and cells defined by issuer, month, and current forbearance status. Because nonbanks’ monthly exercise rate is low in every period, the two benchmarks nearly coincide for them, and the table is informative about the ranking weight rather than about the rationing rule. For nonbanks the actual responses are close to the fixed-ranking computation except in the release, when the response and the latent weight both rise; for depositories the responses exceed the computation from 2021, and the weight rises in 2022. The index for a loan is its index at vesting, so for loans that vested years earlier the spread is measured against an older mortgage rate.

Table IA.5. The Fixed-Ranking Test on the Monthly Eligible Stock

	Pre-shock	Freeze	Release	Recovery		Rate rise
	3/19–2/20	3/20–6/20	7/20–12/20	1/21–6/21	7/21–12/21	1/22–9/22
<i>Nonbanks</i>						
Monthly exercise rate (%)	7.6	2.9	3.3	6.2	7.0	2.3
Relative response	0.80	0.57	1.01	0.71	0.57	0.28
Fixed ranking	0.80	0.62	0.76	0.61	0.61	0.35
Random rationing	0.79	0.61	0.76	0.60	0.60	0.34
Change in latent weight	–	–0.20 (0.52)	1.27 (0.50)	0.48 (0.38)	–0.12 (0.12)	–0.49 (0.12)
<i>Depositories</i>						
Monthly exercise rate (%)	39.1	62.3	36.9	20.0	20.9	16.0
Relative response	0.08	0.00	0.01	0.16	0.12	0.37
Fixed ranking	0.08	0.02	0.03	0.06	0.06	0.13
Random rationing	0.07	0.02	0.03	0.05	0.06	0.07
Change in latent weight	–	–0.36 (0.32)	–0.22 (0.36)	0.65 (0.47)	0.35 (0.49)	0.94 (0.30)

Every loan-month in which a loan is at least three payments delinquent and still in its pool, at depository and nonbank issuers. The ranking index is the loan’s index at vesting. Cells are issuer by month by current forbearance status; the latent weight is estimated with cell fixed effects and period interactions, with standard errors clustered by issuer. Other definitions are as in Table 3.

IA.4.8 Limitations

The index measures the ranking on hard information, and the fixed-ranking prediction assumes that the part of issuers’ ranking it misses is logistic and stable. The index is

estimated on 2016–2018 decisions and ranks later decisions less well as the loan population and the redelivery rules change. The freeze comparison is dominated by a few large issuers (Table IA.3), and a bootstrap over 83 nonbank issuer clusters may understate uncertainty under such concentration; the freeze alone does not separate the fixed ranking from random rationing. All standard errors condition on the estimated index. The disclosures do not follow loans after they leave a pool, so outcomes after repurchase are not observed. The forbearance flag begins in May 2020. Issuer type classifies funding conditions but does not measure them. The comparisons across periods are not estimates of a causal effect of funding; they test whether the choices of constrained and unconstrained issuers are consistent with the model’s rationing rule and its distinction between liquidity and value shocks, against the alternatives described in Section 6.2.

IA.5 Numerical Verification

The script `code/model.py` generates Figures 1–3 and Tables 1 and 2 and writes a verification report, `model_verification.json`, to the `data` folder. It asserts the following.

1. *Propositions 1 and 2.* Seventeen discretized linear programs with 4,000 posterior types reproduce the closed-form private value to within 1.3×10^{-8} . The funded share rises, and the success rate among funded opportunities and the shadow value of cash fall, monotonically in liquidity.
2. *Proposition 4.* For posteriors distributed $\text{Beta}(1.5, 1)$ with $R = 1$ and $c = 0.35$, the decomposition into ranking and screening value holds to machine precision on a grid of 200,000 quantiles. ΔV is concave and peaks at $m = 0.535 = \Pr(Q > p)$ with value $0.1115 = R \mathbb{E}[Q - p]_+ = \frac{R}{2} \mathbb{E}|Q - p|$, and $\Delta V(1) = \mathbb{E}[c - RQ]_+ = 0.0290$.
3. *Proposition 5.* With every action profitable, the slope of a linear probability regression of intervention on the posterior equals $\mathcal{A}(m)/(R \text{Var}(Q))$ at three funded shares, the level slope peaks at m^* , and the slope relative to the base rate is nonincreasing in the funded

share.

4. *Corollary 1.* The binary example peaks at 0.15 at $m = 0.5$, vanishes at both endpoints, and exceeds an adoption cost of 0.06 exactly on $(0.2, 0.8)$.
5. *Proposition 6.* The identity (7) holds for the two scores of Table 2 and for 200 randomly drawn discrete score distributions, with maximum absolute error 2.6×10^{-10} . A mean-preserving contraction of score B is dominated at every fundable share.
6. *Proposition 7.* With $k = 1$, $L = 0.6$, $K_0 = 0.2$, and score B, the fundable share is 0.286 under hard-information pricing and 0.385 under recognition; private value is 0.114 without the signal, 0.149 with it under hard-information pricing, and 0.269 under recognition.
7. *Proposition 8.* In the example of Section 5, private value rises by 0.15, borrower surplus falls by 0.20, and total surplus falls by 0.05, while a planner using the same signal attains total surplus 0.50.
8. *Proposition 9.* In a linear program with 3,000 opportunities and heterogeneous cash needs, a flat ex post payment leaves cash deployed unchanged and raises the number of actions from 0.343 to 0.372; a payment proportional to cash need leaves the allocation unchanged; with a constant cash need, the payment leaves the number of actions unchanged. Under quadratic external-funding costs, a payment and an injection that produce the same funded share move external borrowing in opposite directions.

These checks verify the propositions' numerical statements and the code that implements them. Separately, the scripts in `code/empirics` rebuild Figures 4 and 5, Tables 3–6, Tables IA.2–IA.5, and every number in Sections 6.1–6.7, including the index AUCs, from the source files.