Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Financing Executable Liquidity: Readiness, Prudential Constraints and Certified Policy Choice

Mohammed Ahnouch

mohammed.ahnouch@univ-paris1.fr

Université Paris 1 Panthéon-Sorbonne

Research Article

Keywords:

Posted Date: November 17th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11332813/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Financing Executable Liquidity: Readiness, Prudential Constraints and Certified Policy Choice

Mohammed Ahnouch*

Abstract

Collateral preparation improves access to liquidity but consumes resources before bank capital is measured. We model this timing as a divisible investment and ask when recognising executable capacity improves welfare after the bank finances it. The marginal investment condition prices physical execution and supervisory recognition against the shadow cost of expenditure. In a restricted financing economy, a capital-headroom threshold then identifies when a positive payment-service gain is overturned by the output lost through lending contraction. A multi-bank potential game extends the comparison to equity issuance, collateral portfolios and runnable funding. June 2025 disclosures for 107 European banks motivate retaining size and opening-equity heterogeneity, which changes the financing response in the numerical benchmark. Independent rational bounds translate equilibrium error into auditable welfare comparisons across five recognition rules, including unresolved rankings. Experiments reaching 1,000 banks show how decision precision depends on economic curvature and verification effort. The resulting policy assessment values executable liquidity together with the productive assets and funding adjustments that finance it.

Keywords: executable liquidity; collateral preparation; bank capital; potential games; verified computation.
JEL: C63, C72, G21, G28.

1 Introduction

A bank can hold eligible collateral and still be unable to borrow against it by a payment deadline. Legal documentation, data validation and operational arrangements determine how much of the asset pool can actually be mobilised. The Bank of England's pre-positioning guide describes these tasks and their associated expenses, while Coelho et al. [9] places them within the broader problem of readiness for central bank support [1]. Conditional on facility access, executable capacity is therefore an investment: banks choose how much collateral to prepare and bear its cost before obtaining the liquidity benefit.

That timing creates the central economic trade-off. Preparation expenditure reduces opening cash and equity, so additional readiness may require new capital, a smaller loan book or different funding. A rule that recognises only executable capacity can consequently improve payment service while reducing total surplus. Our central question is: *when does tighter recognition of executable collateral increase total welfare relative to nominal recognition, after banks finance preparation and adjust their opening books?* Answering it requires both an investment condition and an account of

*PRISM — Pôle de recherches interdisciplinaires en sciences du management (UR 4101), Université Paris 1 Panthéon-Sorbonne, 17 rue de la Sorbonne, 75005 Paris, France. Correspondence: mohammed.ahnouch@univ-paris1.fr. ORCID: 0009-0002-6324-8348.

the balance-sheet response. A reduction in delayed payments alone does not price the resources used to obtain it.

The comparison has a conditional answer. Preparation improves welfare when its payment-service value exceeds both its own cost and the productive surplus displaced by financing it. The restricted economy makes this condition an explicit headroom threshold. In the full model, the same strict rule produces certified gains at some cost/equity pairs and certified losses at others. Within the five-rule menu, the reference designs retain several potentially best rules; the disclosure benchmark certifies each of the four less-strict rules above full strictness. These conclusions compare complete equilibria with nominal recognition, rather than comparing payment outcomes at a fixed book.

The financing side of this question follows Froot and Stein [12], who link capital budgeting and risk management when external finance is costly, and Rampini and Viswanathan [15] and Rampini et al. [16], who show how collateral and net worth shape those choices. We add an operational margin that changes the executability of collateral through expenditure incurred before capital measurement. Proposition 1 expresses its marginal value as the sum of physical execution and supervisory recognition benefits, evaluated against financing-adjusted preparation cost. Proposition 3 then isolates a capital-headroom threshold: below it, the output lost through lending contraction outweighs an investment's positive fixed-book value. This threshold gives the financing mechanism an explicit economic boundary.

Recovering that boundary in a richer banking system requires allowing the book to adjust. Strategic responses under stress are central to Hałaj and Priazhkina [13], and Sundaresan and Xiao [17] analyse how liquidity regulation affects lending and liquidity premia. Our multi-bank economy accordingly lets recognition change lending, securities, issuance and runnable funding together. It reconciles shareholders and counterparties so that the welfare comparison measures resources rather than gross financial payments. Section 7.2 shows why this distinction matters: losses can arise through a binding capital floor or through funding contraction while capital remains slack. The two cases identify different financing margins for treasury and supervision to evaluate.

The observed heterogeneity documented in Section 2 motivates retaining bank size and opening equity in that comparison. Section 3 formalises the resulting investment problem; Sections 4 and 5 derive its incentive and welfare implications. Potential games, convex duality and sensitivity analysis [6, 7, 14] make these implications computable, while Section 6 converts numerical equilibrium error into welfare bounds. The experiments in Section 7 use those bounds to compare five recognition rules, including a 107-bank disclosure benchmark and systems of up to 1,000 banks. This emphasis on reproducible policy comparisons complements robust financial-intervention analysis using public EBA data [8]. Throughout, facility recognition is a deadline assessment distinct from Basel's operational requirements for HQLA [2]; this separation keeps the investment incentive and the liquidity-coverage restriction identifiable in the model.

2 Balance-sheet heterogeneity and the readiness margin

The financing question has two distinct empirical dimensions: the resources available to an individual bank and the weight of that bank in the system. We examine both using the EBA's June 2025 reference disclosures at the highest consolidation level [11]. Among 119 named institutions, 107 report the FINREP asset and total-equity observations required for this comparison; the original metadata identifies the remaining twelve as non-FINREP reporters. The complete sample contains EUR 28.28 trillion of assets and EUR 1.94 trillion of accounting equity.

Size is sufficiently concentrated that a representative bank would obscure the relevant welfare weights. Median assets are EUR 96.08 billion, with an interquartile range of EUR 47.74–254.81

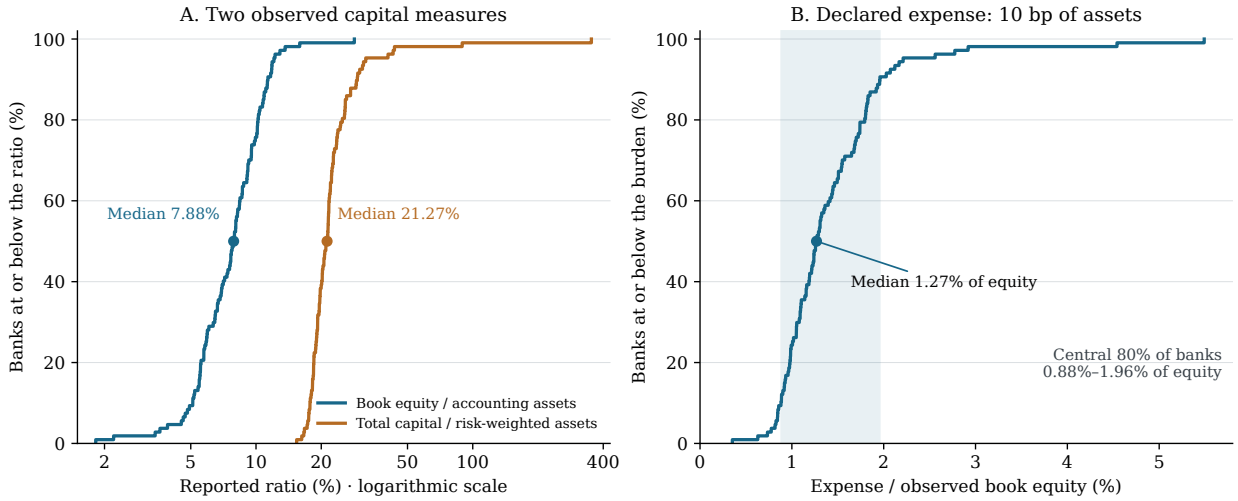


Figure 1: Observed financing heterogeneity among 107 European banking groups, June 2025. Panel A: equal-bank empirical distributions of accounting equity relative to assets and transitional regulatory total capital relative to risk-weighted assets; the logarithmic axis retains every observation. Panel B: an illustrative expenditure of 10 bp of assets expressed as a share of observed accounting equity, with its central 80% interval shaded. Source: EBA transparency disclosures; author's calculations.

billion, while the ten largest groups hold 51.22% of sample assets. Accounting equity relative to assets has a median of 7.88% and an interquartile range of 5.90–9.88%. Its equal-bank mean is 8.13%, compared with an asset-weighted mean of 6.87%. The ten largest groups have pooled accounting equity of 6.21% of assets, compared with 7.56% for the remaining 97. This size composition explains why a common expenditure per asset absorbs a different share of aggregate equity than an equal-bank comparison suggests. These are precisely the two objects the model retains: bank-specific opening equity determines financing choices, and asset weights determine their contribution to system welfare.

Figure 1 relates this dispersion to preparation. Panel (A) compares the accounting-equity distribution with the reported transitional total-capital ratio (median 21.27%). The measures have different capital definitions and exposure denominators: accounting equity is scaled by total assets, whereas the prudential ratio uses regulatory capital and risk-weighted exposures. They therefore answer different financing questions. Accounting equity anchors the opening books in Section 3; the model's capital requirements and asset choices determine the headroom entering Proposition 3.

Panel (B) expresses the same dispersion in the units of an investment budget. For an illustrative resource expenditure of 10 bp of assets, the median expenditure-to-equity ratio is 1.27%, and its central 80% spans 0.88–1.96%. Thus a common asset-scaled budget absorbs more than twice as much equity at the upper end of that interval as at the lower end. This calculation translates the observed balance sheets into financing exposure. Whether that exposure induces issuance, lending contraction or a change in funding is the equilibrium question developed next. The later disclosure benchmark in Section 7.3 holds preparation technology and stress assumptions fixed while introducing these observed sizes and opening-equity ratios; Appendix C records the complementary liquidity disclosures.

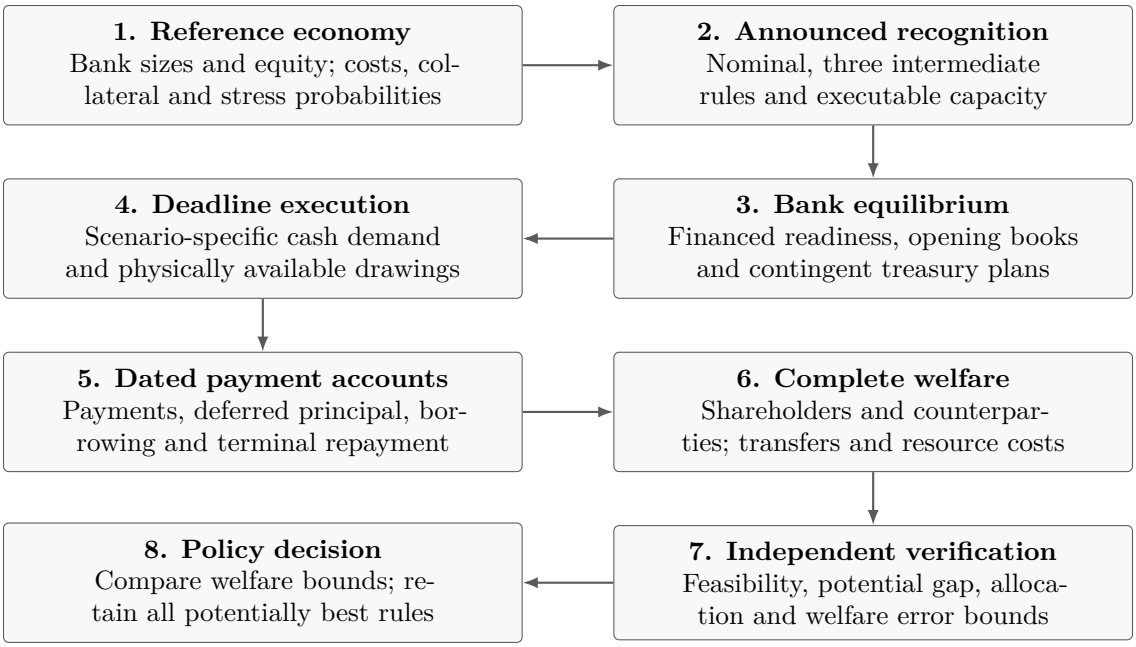


Figure 2: From bank primitives to a certified recognition decision. Each announced rule requires its own equilibrium and welfare account. Treasury outputs are preparation budgets and funding plans; supervisory outputs are the admissible policy comparisons.

3 A divisible readiness economy

The evidence in Section 2 identifies the balance-sheet dimensions along which preparation must be financed. The model now makes their interaction explicit: expenditure first reduces opening resources, physical execution determines the liquidity delivered in stress, and prudential restrictions govern how the bank can respond. Figure 2 follows this sequence through equilibrium, complete accounting and policy verification.

3.1 Dates, quantities and accounts

There are N banks, S jointly observed scenarios and C collateral classes. Bank i has reference asset mass $w_i > 0$, with $\sum_i w_i = 1$; its lower-case quantities are measured per unit of its own reference assets, and aggregates use these weights. Reported numerical amounts are in bp of fixed reference system assets over the model period. Scenario probabilities $p_s > 0$ sum to one. At date 0 the supervisor announces recognition strictness $\rho \in [0, 1]$, after which banks choose opening books, audited readiness and contingent treasury plans. The scenario is revealed at date 1, and principal and contractual compensation are repaid at date 2.

A bank chooses loans l_i, runnable funding b_i, new equity I_i, securities k_{ic}, enabled collateral capacity u_{ic} and preparation spending v_i. Stable funding f_i and inherited equity E_i are given. Reserved capacity is $R_{ic} = h_c k_{ic}$, where $0 < h_c \leq 1$. Preparation has resource cost

$$P_i(u_i) = \sum_c \left(a_c u_{ic} + \tfrac{1}{2} \beta_c u_{ic}^2 \right). \tag{1}$$

Spending and capacity satisfy

$$v_i \geq P_i(u_i) \qquad \text{(preparation spending)}, \tag{2}$$
$$0 \leq u_{ic} \leq R_{ic} \qquad \text{(auditable capacity)}. \tag{3}$$

The epigraph records real expenditure paid before the opening capital measurement. The opening accounts are

$$e_i = E_i + I_i - v_i \qquad \text{(opening equity)}, \tag{4}$$
$$A_i = e_i + f_i + b_i \qquad \text{(balance-sheet funding)},$$
$$m_i = A_i - l_i - \sum_c k_{ic} \geq 0 \qquad \text{(cash availability)}.$$

The distinction between spending and issuance in (4) is essential: preparation reduces measured equity, whereas I_i is the net contribution received by the bank. Shareholders additionally bear issuance resource cost $a_I I_i + h_I I_i^2/2$. Thus financing can replenish capital, but doing so consumes resources that must enter the policy comparison.

The opening book determines the cash available at the payment deadline. Actual facility drawings q_{isc} and contractually deferred principal d_{is} then satisfy

$$0 \leq q_{isc} \leq \xi_{isc}(R_{ic} - u_{ic}) + u_{ic} \qquad \text{(physical execution)}, \tag{5}$$
$$m_i + \sum_c q_{isc} + d_{is} \geq \omega_s f_i + b_i + o_s A_i \qquad \text{(payment-deadline cash)}, \tag{6}$$
$$0 \leq d_{is} \leq \omega_s f_i + b_i \qquad \text{(contractual deferral)}. \tag{7}$$

The execution coefficient $\xi_{isc} \in [0,1]$ determines how much unprepared collateral can be mobilised; preparation replaces that uncertain capacity with executable units. The same scenario also generates operating expenditure $o_s A_i$, which enters both deadline cash demand and terminal profit. Deferral bridges any remaining cash shortage at the cost of compensation $T(d) = \gamma_d d + h_d d^2/2$ at maturity. This transfer reimburses creditors for an equal timing loss, while downstream disruption δd remains uncompensated. Section 3.3 uses this distinction to count each resource cost once.

Securities are newly originated pass-through investments with deterministic net production $r_k k_{ic}$. Holding and origination consume $h_k k_{ic}^2/2$ in resources, and pledged securities remain on the bank's balance sheet. Drawings are repaid with a fee $r_q q$ received by the facility. Their real execution expense is $h_q q^2/2$. Since transactions finance new investments and collateralised borrowing, the securities-buyer account is zero.

3.2 Prudential restrictions and recognition

Preparation and borrowing would be a purely operational choice if the bank could freely replace the resources they use. Prudential restrictions limit that replacement. In addition to the quantity bounds in Appendix B, banks obey the following local rows:

$$e_i \geq \kappa_i(.65 l_i + .20 \textstyle\sum_c k_{ic} + .01) \qquad \text{(internal-risk capital)},$$
$$e_i \geq .725 \kappa_i(l_i + .30 \textstyle\sum_c k_{ic} + .01) \qquad \text{(capital output floor)}, \tag{8}$$
$$e_i \geq .04 A_i \qquad \text{(leverage capital)}, \tag{9}$$
$$m_i + .85 \textstyle\sum_c k_{ic} \geq .10 f_i + .40 b_i \qquad \text{(HQLA coverage)}, \tag{10}$$
$$m_i \geq .06 f_i + .24 b_i \qquad \text{(Level 2 concentration)},$$
$$e_i + .95 f_i + .50 b_i \geq .65 l_i + .15 \textstyle\sum_c k_{ic} \qquad \text{(stable funding)}. \tag{11}$$

These rows represent capital, leverage, LCR and NSFR restrictions [2–5]. The .725 output-floor factor defines the fully phased requirement used in the stress design. The second liquidity row implements the Level 2 concentration cap. HQLA consists of the cash and eligible securities in these two rows.

These liquidity rows govern the opening HQLA stock. Recognition addresses a different object: the cash obtainable from reserved facility capacity by the deadline. Writing $\bar{\xi}_{ic} = \min_s \xi_{isc}$, define assessed capacity

$$K_i^\rho = \sum_c \{[1 - \rho + \rho\bar{\xi}_{ic}]R_{ic} + \rho(1 - \bar{\xi}_{ic})u_{ic}\}. \tag{12}$$

The deadline assessment is

$$m_i + K_i^\rho \geq \omega_S f_i + b_i + o_S A_i \qquad \text{(supervisory recognition)}. \tag{13}$$

Nominal recognition, $\rho = 0$, credits all reserved capacity; strict recognition, $\rho = 1$, credits its worst-state executable amount. Consequently, readiness can create private value in two places: it relaxes physical execution in (5) and the assessment in (13). An intermediate rule changes investment incentives when the latter row binds. Proposition 1 prices these two channels separately.

We also impose a conservative, bank-local sufficient terminal-solvency envelope. It bounds adverse loan margins and funding coupons over the entire declared strategy box, includes operating and execution costs and bounds quadratic delay compensation from above. This envelope may exclude solvent strategies, but ensures full repayment for every admissible unilateral deviation. Bank-local solvency bounds preserve product feasibility while covering rival-dependent market prices.

3.3 Markets and the complete welfare account

The constraints determine which funding responses are feasible; market prices determine their private returns. Define aggregate loans $L = \sum_i w_i l_i$ and runnable funding $B = \sum_i w_i b_i$. The loan margin is $a_i - g_l L$, and the funding coupon is $.008 + g_b B$. The loan-demand slope g_l measures diminishing marginal borrower output and the funding-supply slope g_b measures increasing marginal supplier opportunity cost. A bank internalises its own effect on both prices. Its shareholders maximise $w_i J_i$, where

$$\begin{aligned} J_i = {} & (a_i - g_l L)l_i - \tfrac{1}{2}h_l l_i^2 + \sum_c (r_k k_{ic} - \tfrac{1}{2}h_k k_{ic}^2) - (.008 + g_b B)b_i - \tfrac{1}{2}h_b b_i^2 - .003 f_i \\ & - v_i - a_I I_i - \tfrac{1}{2}h_I I_i^2 - \sum_s p_s \left[o_s A_i + \sum_c (r_q q_{isc} + \tfrac{1}{2}h_q q_{isc}^2) + T(d_{is}) \right]. \end{aligned} \tag{14}$$

Inherited wealth and net new equity principal cancel when evaluating shareholder gain; the real issuance friction remains. The h_l and h_b terms follow the resource schedules in Table 1.

The market-related squares have a separate origin from operational curvature. Borrower output is $\sum_i w_i a_i l_i - g_l L^2/2$; subtracting loan payments leaves borrower surplus $g_l L^2/2$. Funding suppliers incur opportunity cost $.008B + g_b B^2/2$; subtracting this from their receipts leaves supplier surplus $g_b B^2/2$. These two terms therefore follow by integrating the stated demand and supply schedules. They are financial-surplus corrections when aggregating shareholders and counterparties. The facility receives $r_q \sum_{isc} w_i p_s q_{isc}$. Delay compensation and creditor timing loss cancel within the creditor account. The recognition instrument has a zero fiscal account. Consequently,

$$\mathcal{W} = \sum_i w_i J_i + \tfrac{1}{2}g_l L^2 + \tfrac{1}{2}g_b B^2 + r_q \sum_{isc} w_i p_s q_{isc} - \delta \sum_{is} w_i p_s d_{is}. \tag{15}$$

Equation (15) is therefore equivalent to borrower net output and securities production less all lending, funding, preparation, issuance, operating, execution and disruption resources. This identity is the bridge between the investment condition in Section 4 and the welfare threshold in Section 5: prices influence behaviour, while the enlarged account determines whether the resulting allocation creates surplus. Direct cashflow calculations independently reconcile the shareholder and resource evaluations.

3.4 Economic meaning and origin of cost curvature

The welfare account requires an economic interpretation of its cost terms. Each quadratic resource term integrates an explicitly specified affine marginal requirement. For activity z, assume that the next unit uses resources worth $m_z(t) = \alpha_z + \beta_z t$ at scale t, with $\beta_z \geq 0$. Integrating gives

$$C_z(z) = \int_0^z m_z(t)\, dt = \alpha_z z + \tfrac{1}{2}\beta_z z^2. \tag{16}$$

Thus curvature represents increasing marginal workload or adjustment costs. For readiness, a_c prices the first prepared unit and β_c measures the incremental difficulty of enabling further units: scarce legal, data and control capacity, or the inclusion of increasingly complex assets. The Bank of England documents these activities, their costs and variation with collateral complexity [1]. The guide supplies the institutional basis for preparation work; its affine marginal schedule and coefficients define our experimental technology. Fixed costs of initial facility access are sunk in this conditional investment problem.

Table 1 identifies the remaining schedules. Each isolates a separate resource margin. The issuance friction represents real placement, due-diligence and administrative inputs paid outside the bank by shareholders. Linear–quadratic issuance costs also appear in Corbae and D'Erasmo [10]; their formulation supplies a functional-form precedent. Here the specified friction measures consumed resources, while any dilution or price discount transfers wealth between old and new shareholders. This resource account complements the joint financing and risk-management perspective of Froot and Stein [12]. Bank funding administration and suppliers' opportunity costs belong to different accounts. Execution workload belongs to the borrowing bank, while its facility fee is a transfer to the provider. Creditor compensation reimburses the timing loss specified by T.

These schedules are normalised by bank reference assets, so a common schedule represents comparable marginal workloads at different scales. Appendix B states the corresponding monetary-cost identity and coefficients. This convention allows the size and equity heterogeneity in Section 2 to enter the model without silently changing its operating technology.

Within the declared strategy box, the schedules define the experimental technology exactly. Their positive slopes make additional workload increasingly costly and, as Proposition 2 shows, supply the curvature needed to bound allocation error. Section 7.2 varies the preparation intercept and inherited equity to separate technology from financing; Section 7.4 also examines the zero-curvature limit. The computational property thus follows from the stated cost mechanism, rather than determining its economic interpretation.

3.5 The announced-policy game

For a fixed recognition rule ρ, the players are the N banks. Each simultaneously chooses the opening book, preparation and full contingent treasury strategy

$$x_i = (l_i, b_i, I_i, k_i, u_i, v_i, (q_{is}, d_{is})_{s=1}^{S}). \tag{17}$$

Activity	Marginal schedule	Economic interpretation
Preparation u_c	$a_c + \beta_c u_c$	Legal, data and control work for additional executable collateral.
Lending l	$h_l l$	Incremental screening, monitoring and servicing workload.
Runnable funding b	$h_b b$	Bank-side funding administration and liquidity management.
Equity issuance I	$a_I + h_I I$	Increasing resource requirements for arranging a larger issue.
Securities k_c	$h_k k_c$	Origination, due diligence and portfolio administration.
Execution q_{sc}	$h_q q_{sc}$	Increasing operational workload when drawing against collateral.
Deferred principal d_s	$\gamma_d + h_d d_s$	Creditor adjustment and timing loss as more obligations are deferred.

Table 1: Resource-cost assumptions. Integrating each marginal schedule produces the corresponding linear–quadratic cost. Coefficients and the activity domain are specified in Appendix B.

Primitives, the announced rule and the scenario law are common information. The choices in (17) are made before the common stress realisation, and the contingent drawings and deferrals are committed plans implemented when that state is observed. Banks maximise the expected shareholder gain J_i in (14), evaluating the probabilities when choosing the entire strategy.

Write $X_i(\rho)$ for bank i's feasible set defined by its local accounts, prudential rows, solvency envelope and quantity bounds. These restrictions contain no rival choices, so joint feasibility is the product $\prod_i X_i(\rho)$. Aggregate lending and funding prices couple payoffs. A pure noncooperative Nash equilibrium satisfies

$$w_i J_i(x_i^*, x_{-i}^*) \geq w_i J_i(y_i, x_{-i}^*) \quad \forall y_i \in X_i(\rho) \quad \text{(bank optimality)}. \tag{18}$$

Positive asset weights preserve each bank's best responses. Proposition 2 establishes an exact potential for the system-weighted utilities $w_i J_i$; for the bank-unit profits J_i, its identity is $\Delta_i \Phi = w_i \Delta_i J_i$. Thus the same strategic problem is a weighted potential game in bank-unit profit coordinates [14].

The supervisor announces a rule and compares the resulting equilibria using (15). Section 6 makes this comparison over five policies. Announcement fixes the environment for each bank game; the policy exercise ranks those conditional outcomes without specifying a separate strategic objective and equilibrium for a leader.

4 Readiness incentives and equilibrium

The opening accounts in (4) connect operational expenditure to every restriction that uses capital or cash. The first result makes that connection precise. Optimal spending equals the cost of the prepared capacity, and the multiplier on this expenditure aggregates the financing pressures under which the bank invests.

Proposition 1 (Tight expenditure and the readiness shadow value)**.** *Suppose $o_s < 1$, quantity bounds allow a decrease in v_i, and all local constraints are those above or additional rows nondecreasing in v_i. Every optimum has $v_i = P_i(u_i)$. At a KKT point with interior spending, let α_i be the expenditure*

multiplier, λ_{isc} the execution multipliers and η_i the recognition multiplier, all normalised by w_i. For $0 < u_{ic} < R_{ic}$,

$$\alpha_i(a_c + \beta_c u_{ic}) = \sum_s (1 - \xi_{isc})\lambda_{isc} + \rho(1 - \bar{\xi}_{ic})\eta_i. \tag{19}$$

Moreover $\alpha_i = 1 - \sum_s p_s o_s$ plus the nonnegative expenditure loadings of the binding cash, capital, leverage, liquidity, recognition and solvency rows.

Proof. Reducing excess spending raises opening equity and cash. In a cashflow or recognition row the improvement is multiplied by $1 - o_s > 0$; every other stated row also relaxes. Shareholder profit rises at rate $1 - \sum_s p_s o_s > 0$. Excess spending is therefore dominated. Introduce the multiplier for $P_i(u_i) - v_i \leq 0$. Stationarity in v_i gives the expression for α_i. Stationarity in u_{ic} gives (19); the capacity upper and nonnegativity multipliers vanish for an interior choice. At zero readiness the marginal benefit cannot exceed the marginal financing-adjusted cost, and at full readiness the inequality reverses, with complementary multipliers included. Constraint qualification is needed for this multiplier representation, not for the preceding dominance argument. Multipliers need not be unique at degenerate points. □

Equation (19) supplies a treasury investment rule: prepare another unit when its execution and recognition benefits cover the financing-adjusted marginal cost. A binding capital requirement increases that cost, linking the rule directly to opening equity. But the rule remains a private optimality condition; it does not include every consequence for borrower output and counterparties in (15). To compare those consequences across policies, we first establish how the banks' simultaneous choices determine an equilibrium.

Proposition 2 (Potential, existence and conditional uniqueness)**.** *Suppose each bank's feasible set is nonempty and all quantity bounds are finite. The game has a pure equilibrium. A convex minimisation finds every equilibrium when the negative potential is convex; under the strictly positive curvatures specified below it finds a unique reduced allocation.*

Proof. For the two aggregate markets, subtract from $\sum_i w_i J_i$ the cross interactions counted twice and obtain a potential Φ whose negative, up to an irrelevant constant, is

$$F(x) = \tfrac{1}{2} x^\top D x + c^\top x + \tfrac{1}{2} g_l L^2 + \tfrac{1}{2} g_b B^2. \tag{20}$$

Here the diagonal loan entry is $w_i h_l + g_l w_i^2$, the funding entry is $w_i h_b + g_b w_i^2$, and the other positive quadratic entries are the weighted local resource costs. The spending coefficient is $w_i(1 - \bar{o})$, where $\bar{o} = \sum_s p_s o_s$. For any feasible unilateral change by bank i,

$$\Phi(y_i, x_{-i}) - \Phi(x_i, x_{-i}) = w_i\{J_i(y_i, x_{-i}) - J_i(x_i, x_{-i})\}. \tag{21}$$

This is exact potentiality for the utilities in (18). Product feasibility is crucial. Maximising a continuous potential on this compact set yields an equilibrium. Conversely, in the convex product problem the concatenation of all block optimality inequalities is the global variational inequality, so a Nash equilibrium minimises F.

Eliminate spending using Proposition 1. Every affine row has a nonnegative coefficient on v_i, so substitution preserves convexity of the reduced feasible set. When $h_l, h_b, h_I, h_k, h_q, h_d, \beta_c$ and every p_s, w_i are positive, the reduced objective is strongly convex with the explicit valid lower bound

$$\mu = \min\{D_{jj} : D_{jj} > 0;\ w_i(1 - \bar{o})\beta_c : i, c\} > 0. \tag{22}$$

The two aggregate squares add positive semidefinite curvature. The reduced allocation is unique. If some curvatures vanish, existence persists but this uniqueness and distance certificate need not follow. □

Proposition 2 provides the equilibrium object used in the rest of the paper. Its finite strategy boxes ensure existence and support the global error bounds in Section 6; each experiment verifies them together with a strictly feasible cash-rich allocation. The welfare question can now be stated as a comparison between equilibria, rather than as the marginal value of preparation at a fixed book.

5 An explicit capital-headroom threshold

The private investment condition does not determine whether improved execution survives its financing response. We isolate that response in a restricted economy with fixed cash, elastic long-term funding and a capacity-only recognition assessment. These restrictions leave the capital-funded loan contraction as the sole portfolio margin, producing a closed-form benchmark for the richer adjustments examined in Section 7.2.

One bank has a loan opportunity cap $\bar{l}$, capital ratio $\kappa > 0$ and inherited capital $E = \kappa\bar{l} + h$, where $h \geq 0$ is headroom. Net readiness expenditure is $P(u) = au + \beta u^2/2$. Stable funding adjusts residually to the opening accounts at zero incremental resource cost. A fixed collateral pool supports R units of capacity; its constant cost is omitted from comparisons. Unprepared collateral is unavailable at the deadline, and cash is fixed, so actual drawings equal u and deferred principal equals $D - u$ for $u \leq D$. The supervisory and lending restrictions are

$$\begin{aligned} (1-\rho)R + \rho u &\geq M && \text{(capacity recognition)},\\ \kappa l &\leq E - P(u) && \text{(loan-capital requirement)},\\ 0 \leq l &\leq \bar{l} && \text{(loan opportunity)},\\ 0 \leq u &\leq \min\{R, D\} && \text{(usable readiness)}. \end{aligned}$$

Thus the minimum admissible readiness is

$$u_\rho = \max\{0, R - (R-M)/\rho\} \quad (\rho > 0), \qquad u_0 = 0. \tag{23}$$

The parameter domain is $0 < M < R$ and $u_1 \leq D$, so each recognition rule is feasible. Loan demand is $A - \gamma l$, shareholder revenue is $(A - \gamma l)l$ and borrower output is $F_L(l) = Al - \gamma l^2/2$. Assume revenue is increasing on $[0, \bar{l}]$, the bank's marginal compensated delay benefit is at most t, and $P'(u) \geq t$. The smallest admissible u is then optimal, and is the unique readiness choice if the marginal inequality is strict except at isolated points, as in the printed family. At this optimum the bank chooses the largest capital-feasible loan book:

$$l(u) = \bar{l} - \frac{(P(u) - h)_+}{\kappa}. \tag{24}$$

Indeed, raising u consumes resources and weakly reduces a profitable loan book; its direct marginal cost already exceeds its private payment benefit. Elastic stable funding and a sufficiently large positive funding domain implement the accounts without forcing additional runnable borrowing. All comparisons below require $l(u) \geq 0$.

Let total marginal delay resource cost, including the uncompensated part, be $\Delta > t$, and set $p_0 = F_L'(\bar{l}) > 0$. Define $z = (P(u) - h)_+/\kappa$ and $K = \Delta u - P(u)$, the payment-service gain net of preparation at a fixed loan book.

Proposition 3 (Sharp financing threshold in the restricted economy)**.** *For an equilibrium readiness level u, the full welfare change from $u = 0$ is*

$$\Delta\mathcal{W} = K - p_0 z - \tfrac{1}{2}\gamma z^2. \tag{25}$$

If $K < 0$, readiness lowers welfare for every headroom h. If $K \geq 0$ and $\gamma > 0$, define

$$z_* = \frac{\sqrt{p_0^2 + 2\gamma K} - p_0}{\gamma}, \qquad h_* = P(u) - \kappa z_*. \tag{26}$$

Then welfare is weakly higher exactly when $h \geq \max\{0, h_*\}$. For $\gamma = 0$, replace z_* by K/p_0. Preparation can have $K > 0$ and nonetheless reduce equilibrium welfare for $0 \leq h < h_*$.

Proof. The gain from less delay is Δu, and preparation consumes $P(u)$ once. Loan contraction loses $F_L(\bar{l}) - F_L(\bar{l} - z) = p_0 z + \gamma z^2/2$. This proves (25). The loss is increasing in $z \geq 0$, and z decreases with headroom. Solving its equality with K yields the threshold. The statement for $\gamma = 0$ is its linear limit. □

The proposition gives a direct construction of the financing reversal. Fix a feasible positive readiness choice, set initial headroom to zero, and write $z = P(u)/\kappa$. Then choose the payment-loss parameter in the interval

$$\frac{P(u)}{u} < \Delta < \frac{P(u) + p_0 z + \gamma z^2/2}{u} \qquad \text{(financing-reversal region)}. \tag{27}$$

The left inequality makes the operational investment valuable at a fixed book; the right makes the lost loan output larger than that gain. Choosing the private compensation rate below marginal preparation cost also ensures that the bank prepares only the amount required by recognition. The example therefore follows directly from the two economic inequalities in Proposition 3.

Table 4 chooses simple rational values inside this region. Strict recognition requires two-thirds of the reserved capacity to be ready. Preparation yields a fixed-book gain of 55.56 bp, but financing it from binding capital displaces loan output worth 147.05 bp, leaving an adjusted loss of 91.50 bp. Holding the book fixed requires restoring the capital consumed by preparation; the comparison continues to charge preparation as a real resource cost. It isolates the operational benefit that the equilibrium financing response then overturns.

The threshold also explains the slope of the readiness frontier. Holding u fixed where $K > 0$ and $h_* > 0$ gives

$$\frac{\partial h_*}{\partial a} = u\left(1 + \frac{\kappa}{\sqrt{p_0^2 + 2\gamma K}}\right) > u. \tag{28}$$

A higher preparation price therefore increases required headroom by more than its direct expenditure effect: it also reduces the net payment-service gain available to absorb lending contraction. The proved threshold determines the continuous region, and 1,476 rational cost/headroom evaluations generate its plot. Section 7.2 returns to the multi-bank economy, where additional funding and collateral margins enter the comparison. Before doing so, we establish how its computed policy differences can be certified.

6 Certificates, policy decisions and bank-local computation

The headroom threshold has an exact expression because the restricted economy has one financing margin. In the full model, the same economic comparison requires solving the simultaneous portfolio choices of Section 3. Proposition 2 makes this a convex optimisation problem; the certificate below translates its numerical error into the precision needed for a policy decision.

Let x be an exactly feasible point of the epigraph formulation, $\lambda \geq 0$ multipliers for affine rows and $\nu \geq 0$ multipliers for the preparation epigraphs. Let $\mathcal{L}$ be the convex Lagrangian and $[\ell, r]$ the box containing every feasible allocation. A supporting hyperplane gives

$$F(x) - F(x^*) \leq F(x) - \mathcal{L}(x) - \min_{\substack{\ell \leq y \leq r \\ \text{strategy box}}} \nabla \mathcal{L}(x)^\top (y - x) =: \varepsilon. \tag{29}$$

A nonnegative diagonal curvature minorant H of the Lagrangian sharpens the bound by adding $\sum_j H_{jj}(y_j - x_j)^2/2$ inside the minimisation. Each coordinate minimum has a closed form: project $-\nabla_j \mathcal{L}/H_{jj}$ into its displacement interval, using the endpoint rule when $H_{jj} = 0$. The checker verifies feasibility, multiplier signs and all arithmetic using exact rational representations of the stored binary coefficients. Large rational support terms use directed upward rounding before addition, preserving the upper bound.

The certificate bounds every unilateral improvement in the weighted shareholder payoff $w_i J_i$ by ε, in units of reference system assets. For bank i's profit per unit of its own reference assets, the corresponding bound is ε/w_i. Set $\widetilde{x}$ equal to x with spending replaced exactly by $P(u)$. This preserves feasibility and weakly decreases F. Strong convexity of the reduced objective gives

$$\|\widetilde{x}_{-v} - x^*_{-v}\|_2 \leq \sqrt{2\varepsilon/\mu}. \tag{30}$$

The subscript matters: the epigraph objective is not strongly convex in its spending coordinates. A known global gradient bound L_W on the reduced box yields a welfare error at most $L_W\sqrt{2\varepsilon/\mu}$. The point value and bound constants are evaluated with exact rational arithmetic and outward conversion. Expected-delay bounds follow from its linear coefficient norm. The intervals measure deterministic numerical error conditional on the stated primitives. The reported welfare intervals use a sharper diagonal-metric bound. Let $D_{\mathrm{red}} > 0$ be the local reduced curvature diagonal, and let g_j bound the absolute jth welfare derivative over the reduced box. Weighted Cauchy–Schwarz gives

$$|\mathcal{W}(\widetilde{x}) - \mathcal{W}(x^*)| \leq \sqrt{2\varepsilon \sum_j g_j^2/(D_{\mathrm{red}})_{jj}}. \tag{31}$$

The diagonal metric in (31) retains the bank-size information motivating Section 2, whereas a single minimum-curvature bound treats every coordinate at the weakest curvature. Both bounds remain sensitive to rare scenarios and weak marginal costs. Their practical difference appears in the scaling exercise of Section 7.4.

Nominal recognition, $\rho = 0$, supplies the baseline: it retains the same collateral, facility access, HQLA restrictions and stress law while crediting all reserved capacity. Comparing other rules with this baseline isolates the effect of recognition incentives and the resulting bank adjustment. The evenly spaced menu $\rho \in \{0, 1/4, 1/2, 3/4, 1\}$ includes both economically defined endpoints and three transparent intermediate assessments. For each rule, let $[\underline{W}_\rho, \overline{W}_\rho]$ be the resulting interval. Policy ρ is certified better than ρ' only when $\underline{W}_\rho > \overline{W}_{\rho'}$. The retained set contains policies whose upper bound is at least the greatest lower bound. A singleton is a certified best choice within this menu. Overlapping intervals are reported as unresolved, including economically small differences and slack intermediate rules. The policy decision domain is this five-rule menu.

This decision rule also clarifies the computational task: solve each announced policy accurately enough to distinguish welfare intervals, while retaining unresolved alternatives. Only L and B couple the banks. Holding the two aggregate marginal prices fixed decomposes the convex problem into bank-local subproblems. In the reduced metric induced by their positive local curvature, the

aggregate best-response price map has the conservative contraction bound

$$q \le \max\left\{ g_l \sum_i \frac{w_i^2}{w_i h_l + g_l w_i^2}, g_b \sum_i \frac{w_i^2}{w_i h_b + g_b w_i^2} \right\} < 1 \tag{32}$$

for the printed coefficients. To see this, the local strong-monotonicity inequality bounds the response to a linear price perturbation by its inverse-diagonal norm; premultiplication by the two disjoint aggregate vectors gives the displayed operator bound. Exact subproblem solves therefore give geometric contraction. Approximate solves add inner-solve error and need not attain an arbitrarily small iterate tolerance. Our batched implementation allows 120 iterations and certifies the accuracy of its final allocation; the stopping tolerance and certificate are reported separately.

7 Numerical results and economic mechanisms

The numerical comparisons ask how the financing mechanism in Proposition 3 changes when equity issuance, collateral choice and runnable funding can adjust together. We begin with a transparent reference economy, vary financing pressure, and then introduce the observed heterogeneity from Section 2. Every welfare sign uses the certificate of Section 6.

7.1 Reference design and policy comparisons

A ten-bank economy with three scenarios and two collateral classes makes the individual financing and execution margins inspectable before introducing the disclosure sample. Stable funding is 60% and inherited equity 11% of each bank's reference assets; the effective capital parameter is 16%. Together these values leave scope for both equity issuance and portfolio adjustment. A second family raises issuance costs and reduces inherited equity to examine the capital-boundary mechanism. These are complementary stress designs, distinguished by their financing conditions.

Loan-demand intercepts fluctuate around 13% by 0.6 percentage points using fixed, dispersed bank phases; this introduces reproducible return heterogeneity. Securities produce 10% over the model horizon. The two collateral classes have advance factors 65% and 90%, corresponding to haircuts of 35% and 10%; their different mobilisation values create a portfolio allocation margin. These return, cost and advance-factor specifications compare model mechanisms on a common asset scale. The disclosure benchmark subsequently replaces bank sizes and opening equity with their observed values.

Scenario probabilities are in the ratio $1 : 1.5 : 2$, giving more weight to the more adverse states in this stress exercise. Withdrawal fractions span 5–25% of stable funding and operating expenditure spans 20–100 bp of opening assets. Base preparation intercepts of 0.5%, 2.5% and 5% span substantial voluntary investment, partial investment and almost no preparation under nominal recognition. The baseline issuance intercept of 5% represents the resource cost of arranging additional equity; the constrained family doubles it to 10%. Remaining marginal slopes are specified in Section 3.4 and Appendix B. Their purpose is to expose the relative cost of expanding each activity and permit independently bounded policy comparisons. Weights and probabilities are rounded by largest remainders onto a 2^{-48} grid with exact unit sums.

Figure 3 compares the five recognition rules under the three preparation-cost schedules. With a 0.5% intercept, banks already prepare 2,513.46 bp of capacity under nominal recognition. Strict recognition adds only 11.05 bp and eliminates the remaining 6.44 bp of expected deferred principal. Its welfare point gain is 0.086 bp, smaller than the certified uncertainty, so the ranking remains unresolved. At 2.5%, nominal readiness is 714.68 bp and strictness adds 362.20 bp; the welfare

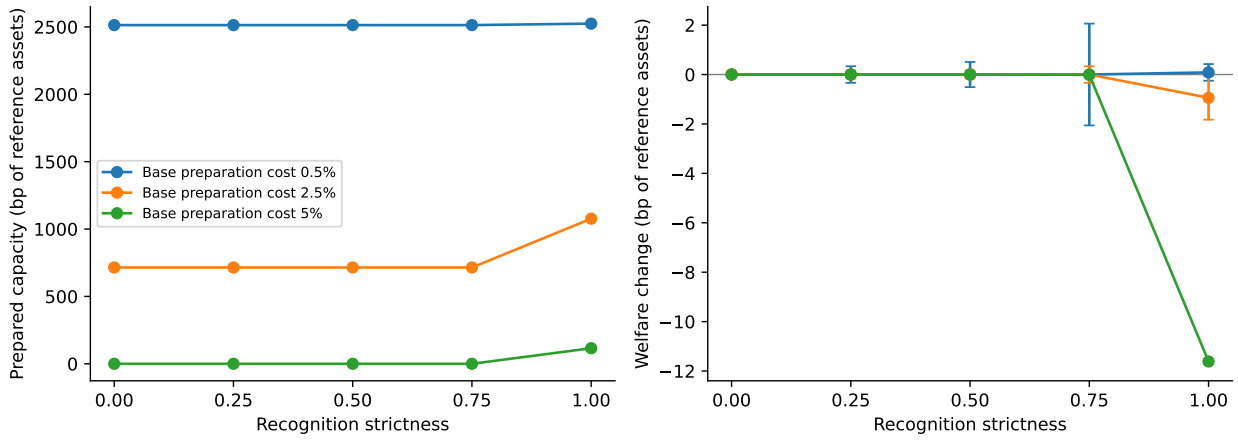


Figure 3: Prepared capacity and welfare changes, in bp of fixed reference system assets, across five announced recognition rules. The cost legend reports the base initial marginal preparation cost a as a percentage of the amount prepared; class-specific intercepts scale it by factors from .85 to 1.15. Lines connect evaluated policies. Error bars subtract the nominal-policy welfare interval from each independently valid policy interval; the baseline compared with itself is exactly zero.

change is −0.94 bp. At 5%, voluntary readiness is negligible, while strictness induces 115.45 bp, eliminates 390.92 bp of expected delay and reduces welfare by 11.61 bp.

The flat segments reflect slack recognition assessments. At the 10^{-6} slack tolerance, all ten recognition rows are slack at 0.5% and 2.5% through $\rho = .75$ and become binding at $\rho = 1$. At 5%, one row already binds at .75 and all ten bind at 1. The plotted final decline therefore compares two menu points on different sets of binding constraints; the connecting segment depicts that comparison. For the expensive preparation technology, banks also reduce runnable funding by 1,122.30 bp and loans by 215.59 bp between nominal and strict recognition. This liability and portfolio response explains why its welfare decline is much larger even though its readiness increase is smaller. More capacity prepared is not, by itself, a measure of policy benefit.

The low-cost welfare interval at $\rho = .75$ is visibly wider because that allocation has a larger certified numerical error bound. The error bars measure computational precision. Under the retained menu, all five rules remain possible best choices at 0.5%; at 2.5% and 5%, strict recognition is eliminated and the four remaining rules have overlapping welfare bounds. This set-valued result directs additional computational precision toward close comparisons and distinguishes them from the resolved losses.

7.2 From execution gains to financing adjustments

The reference comparison shows that a large payment-service gain can coexist with a welfare loss. Proposition 3 predicts one reason: preparation crowds out output when capital is costly to replace. To expose that margin, the constrained-financing family doubles the linear issuance resource cost to 10% and varies inherited equity and preparation cost. Binding-row diagnostics show when the output floor, liquidity cap, recognition row or conservative solvency envelope determines choices. Table 2 reports selected full-model point comparisons. The $a = 4\%, E = 2\%$ economy has all ten output-floor rows binding under strict recognition, while payment service improves and welfare falls. This economy links the loss to capital-financed preparation; the richer model also identifies additional channels through funding and asset composition.

Design	ΔW (bp)	Δd (bp)	Δu (bp)	Sign
$a = 1.5\%, E = 2\%$	0.54	-65.44	152.19	positive
$a = 2.5\%, E = 3.5\%$	-1.58	-217.83	435.22	negative
$a = 4\%, E = 2\%$	-6.79	-309.26	428.29	negative
$a = 6\%, E = 11\%$	-18.82	-462.45	82.29	negative

Table 2: Strict minus nominal recognition in selected full-model economies. All changes are bp of fixed reference system assets; d is expected deferred principal and u is prepared capacity. The design reports preparation cost a and opening equity E as percentages of their respective quantity and reference-asset denominators. Signs use the independent welfare intervals.

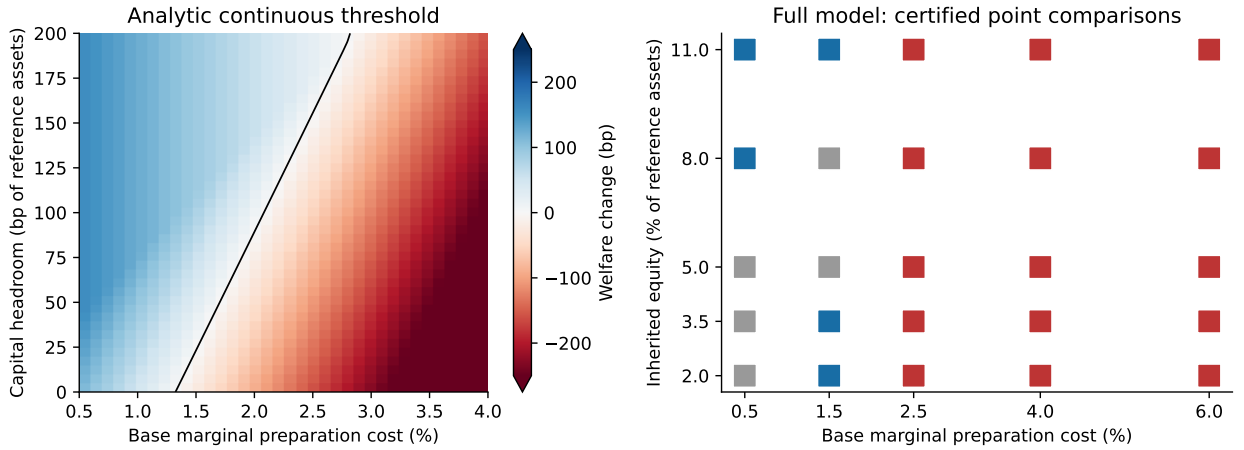


Figure 4: Left: the proved continuous threshold in the restricted financing economy; the curve marks zero adjusted gain. Right: 25 independently checked full-model comparisons at printed parameter points. Blue is a certified gain, red a certified loss and grey an unresolved comparison. The full-model map contains five certified gains, fifteen certified losses and five unresolved comparisons. Preparation prices are percentages of the amount prepared. Headroom and welfare use bp of reference assets; the multi-bank vertical axis gives inherited equity as a percentage of reference assets.

Figure 5 distinguishes two economic mechanisms. At $a = 4\%, E = 2\%$, strict recognition raises preparation expenditure by 16.77 bp while reducing new equity by 291.55 bp, loans by 195.21 bp and runnable funding by 634.45 bp. All ten output floors bind after adjustment. With costly issuance, banks meet the more demanding execution assessment partly by shrinking the book and its payment obligations. Reduced timing and downstream losses contribute 37.85 bp to welfare, and funding and issuance savings contribute 41.58 bp; lower borrower and securities output subtract 79.02 bp, alongside the preparation bill. Remaining lending, execution and operating savings add 9.59 bp. The reconciled net effect is a loss of 6.79 bp. Thus additional readiness can coexist with less newly raised equity: equity issuance and investment respond jointly.

At $a = 6\%, E = 11\%$, a different financing margin becomes decisive. Capital requirements are slack under both policies, while runnable funding falls by 1,386.47 bp, loans by 314.88 bp and securities by 727.45 bp; readiness rises only 82.29 bp. Banks principally reduce the liabilities and assets generating the deadline demand. The resulting welfare loss of 18.82 bp shows why capital headroom alone cannot select the policy in the full model: recognition interacts with funding demand and liquidity composition even away from a capital boundary. The small preparation-cost bars

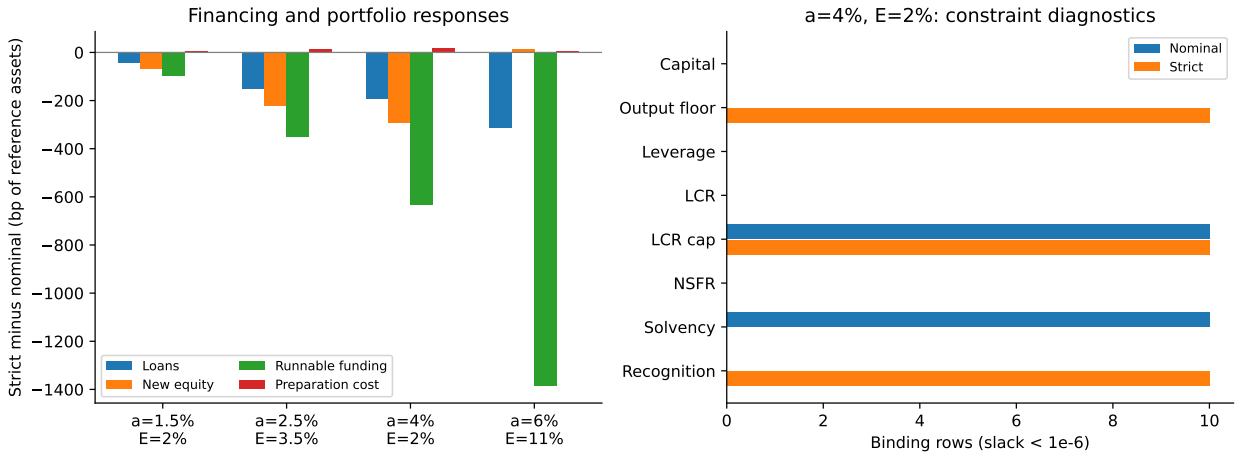


Figure 5: Left: financing and portfolio adjustments in bp of fixed reference system assets for the four economies of Table 2. Right: counts of prudential, solvency and recognition rows with numerical slack below 10^{-6} in the selected constrained economy. Solvency has one row per scenario; other displayed categories have one row per bank. Binding status uses a numerical slack tolerance; welfare signs use the independent intervals.

relative to the asset-adjustment bars display this amplification directly: the cost of investing in readiness and the size of the induced book response are different economic quantities.

Lower preparation costs can reverse this ranking. At $a = 1.5\%, E = 2\%$, the preparation price places the economy on the other side of the trade-off: eliminating 65.44 bp of expected deferred principal saves 7.41 bp in timing and downstream losses. Funding, issuance and operating savings add about 11.11 bp, and lending and execution savings add 0.67 bp. Together they cover the 3.93 bp preparation increase and the 14.72 bp reduction in borrower and securities output. Welfare rises by 0.54 bp. Taken together, the point comparisons identify the relevant supervisory diagnostic: compare the avoided payment loss with preparation expenditure and the net resource effect of the entire financing response.

7.3 Recognition with observed size and equity heterogeneity

The financing cases above hold initial bank size and equity fixed within each design. Section 2 shows why that simplification matters: the disclosure sample contains large differences in both, so counting banks and weighting their assets describe different economic objects. We now set w_i from observed assets and E_i from observed accounting equity relative to assets for the 107 complete observations, retaining the specified preparation technology and stress law. The exercise thus asks how the recognition comparison changes when the opening balance-sheet distribution is replaced by the observed one.

All five policies are recomputed. In this benchmark, all four less-strict rules are certified to exceed the welfare of full strictness, but their mutual ranking remains unresolved. The financing mix explains the resulting comparison. Relative to nominal recognition, strictness raises new equity by 56.33 bp while reducing runnable funding by 548.11 bp and loans by 118.58 bp. Prepared capacity rises 384.66 bp and expected delay falls 293.90 bp. Timing and downstream savings of 35.80 bp are offset by preparation and issuance costs of 17.11 bp and reduced productive output; the total welfare difference is −1.62 bp. The aggregate response combines additional issuance and book contraction across banks with different opening-equity ratios. It therefore prices readiness

N	Variables	Solve (s)	Verify (s)	Gap (bp)	μ
10	170	0.007	0.015	5.62e-06	5.55e-04
50	850	0.020	0.074	5.63e-07	1.11e-04
200	3400	0.062	0.294	5.65e-07	2.77e-05
1000	17000	0.329	1.647	5.72e-08	5.55e-06

Table 3: Bank-count sweep at fixed system size. Potential gaps are bp of fixed reference system assets; solve and verification times are seconds. The conditioning bound μ is rounded downward in the original reduced coordinates; gap bounds are rounded upward. Appendix B states the reporting-unit conversion.

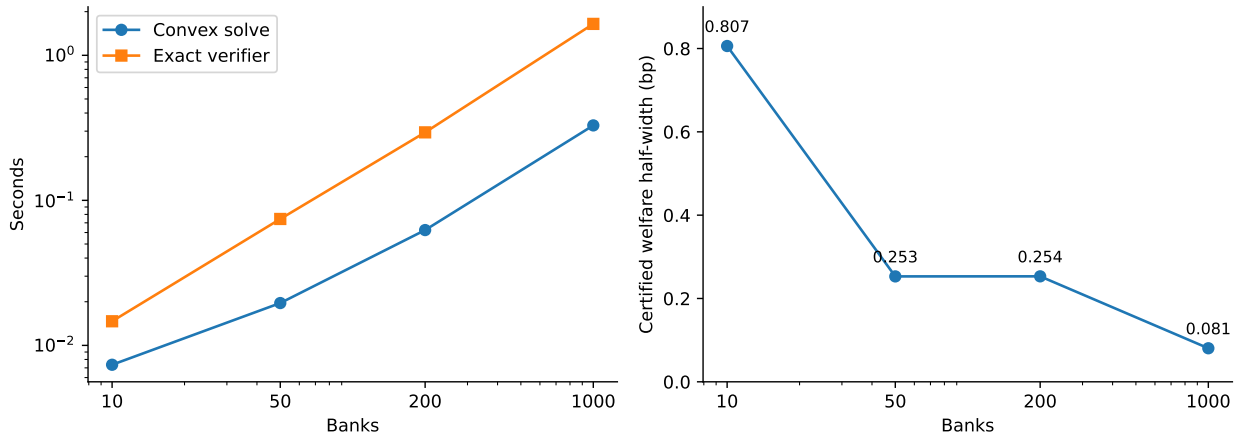


Figure 6: Solve and verification times and certified welfare half-widths. Each half-width is one half of the equilibrium welfare interval, in bp of fixed reference system assets. Points use the computed widths; numerical labels round their exact bounds upward to three decimal places. Timings describe one process and one machine; recorded memory is the process high-water mark.

across a distribution of financing positions, extending the common-equity designs of Section 7.2. The descriptive distributions motivate this exercise; its welfare magnitudes belong to the specified equilibrium model.

7.4 Decision precision as the economy grows

A policy comparison that depends on a fraction of a basis point requires more than a quickly computed allocation. To assess the precision available as the economy grows, Table 3 reports the bank-count sweep at fixed scenario and collateral dimensions. Total reference mass remains one, so the bank-count sweep splits the same system assets across more banks. This isolates organisational dimension from an increase in system size. Bank masses, scenario probabilities and class-specific cost curvatures enter μ explicitly; increasing dimension does not preserve conditioning automatically. The code also varies S through 1, 6, 12 and 24 at $N = 50$, and C through 1, 4 and 8. A combined case uses 200 banks, 12 scenarios and four classes. The collateral-class sweep varies both the feasible asset mix and per-class costs, allowing computational demand to be assessed across richer technologies.

The execution scenarios mix a common adverse component with deterministic bank-specific scenario variation. The common-stress weight determines the mixture; realised execution correlation

is calculated from its joint law. Every experiment reports the actual probability-weighted average pairwise execution correlation. Separate runs use weights 0, .5 and 1. The joint laws are experimental assumptions and vary both marginal execution distributions and dependence. A useful structural implication follows from separability: if each bank's marginal scenario law and feasible contingent plans remain fixed, changing only cross-bank coupling leaves expected payoffs and the equilibrium unchanged. The objective sums bank-local expectations, and aggregate-market quantities are chosen before the scenario is revealed. Hence readiness under this expected-surplus criterion responds to marginal execution reliability. A dependence-sensitive system-risk objective would require an additional aggregate risk criterion.

Across 86 final runs, 85 receive exact convex-program gap and positive-curvature distance bounds; the zero-preparation-cost and zero-curvature case receives a gap-only check. An optimizer-free pass verifies all 87 exported fixtures, including that case and the decomposition output. The latter has contraction bound .75145 and a final certified potential gap 5.64×10^{-7} bp after 120 iterations. Its allocation differs from the central solution by 2.93×10^{-7} in Euclidean norm. Inner-solve variation prevents its requested iterate tolerance from being reached, so the certificate supplies the decision-relevant accuracy measure.

The bank-count sweep increases decision variables from 170 to 17,000. Solve time rises from .007 to .329 seconds and independent verification from .015 to 1.647 seconds on the recorded machine. At 1,000 banks, verification uses about 83% of the combined solve-and-check time. The separation between the runtime curves thus measures the cost of obtaining an independently auditable decision in addition to a numerical allocation.

The accuracy of that decision depends on curvature as well as computational effort. Over the same sweep the Euclidean curvature bound falls from 5.55×10^{-4} to 5.55×10^{-6} because total system mass is held at one. This decline explains why scalability must be judged by certified decision precision as well as raw runtime: a smaller potential gap is converted through the actual economic curvature.

Figure 6 expresses the resulting accuracy in the units of the policy decision. Its welfare half-width bounds are 0.807, 0.253, 0.254 and 0.081 bp at 10, 50, 200 and 1,000 banks, respectively. The decreases accompany smaller feasibility-repair mixtures: 10^{-8}, 10^{-9}, 10^{-9} and 10^{-10}, respectively. These mixtures move each numerical solution toward a verified feasible allocation and contribute to the final error bound. Thus the downward steps describe achieved numerical precision, while the 50–200 plateau reflects similar certificate accuracy. Increasing bank count itself does not improve the economic approximation. Equation (31) retains bank-size information in this conversion. Together with the actual curvature and the recorded feasibility repairs, it explains how a computational gap translates into an economically interpretable policy interval.

8 Treasury and supervisory interpretation

Relative to nominal recognition, stricter assessment improves welfare when the payment loss avoided exceeds preparation expenditure and the net resource cost of the induced financing response. Proposition 3 states this condition exactly in the restricted economy: for positive fixed-book value, sufficiently high headroom preserves the gain, while lower headroom can reverse it. The full model adds issuance, collateral and funding adjustment, producing both certified gains and certified losses at the tested cost/equity pairs.

The five-rule comparison translates this conditional answer into a supervisory choice. In the reference family, all rules remain potentially best at the lowest preparation cost; at the two higher costs, full strictness is eliminated while the remaining four rules overlap. In the disclosure benchmark,

every less-strict rule has certified higher welfare than full strictness, and their mutual ranking is unresolved. Hence service improvement supports a welfare recommendation only after the financing comparison and its numerical bounds have been evaluated.

For treasury, (19) makes the first stage operational. A bank can compare the marginal readiness benefit with the financing cost of another prepared unit, identify which collateral classes justify expenditure, and inspect whether capital, liquidity coverage or operational recognition limits the chosen plan. The model requires that the planned capacity is auditable and that funding compensation is enforceable. These arrangements give operational meaning to the optimised capacity.

For supervision, the second stage requires the complete account in (15). Comparing an announced recognition menu using full balance-sheet responses allows the supervisor to distinguish reductions in payment delay from increases in welfare and retain several policies when uncertainty bounds overlap. The threshold result shows why initial headroom matters: the same useful operational investment can crowd out valuable lending after capital becomes scarce. The full-model point map identifies where operational savings outweigh financing adjustments and where book contraction reverses that gain. Its equity dimension also permits collateral, issuance and funding adjustment, whereas the analytic frontier fixes those additional margins; the exact monotone headroom threshold therefore pertains to the restricted financing economy. Grey full-model points mark unresolved numerical comparisons.

The results apply to banks with existing facility access, auditable readiness and enforceable repayment. The terminal-solvency envelope supplies a conservative feasible domain; the cost schedules and stress laws determine the marginal trade-offs within it. A treasury implementation can estimate those schedules from preparation workloads, issuance expenses and execution records, then rerun the policy menu with its own capital headroom. For supervisors, the central implication is to assess readiness expenditure together with the funding and productive assets that finance it. The numerical certificates establish the precision of that comparison for each specified economy.

A Construction of the rational financing example

The parameters in Table 4 are chosen to satisfy three conditions: recognition induces a positive, interior readiness level; preparation has positive fixed-book social value; and capital-funded loan contraction more than absorbs that value. At zero headroom, the choice $p_0 = \kappa$ makes the first-order output loss equal to the preparation bill. Positive demand curvature supplies the additional loss from moving along the loan-demand schedule.

For these values, the interval in (27) is $13/600 < \Delta < 18889/432000$, and $\Delta = 3/100$ lies in its interior. Strict recognition gives $u = 2/3$, preparation cost $P = 13/900$, loan contraction $z = 13/180$ and remaining loans $l = 257/180$. The fixed-book gain is $K = 1/180$, while foregone loan output is $p_0 z + \gamma z^2/2 = 9529/648000$. Hence the adjusted welfare change is $-5929/648000$. Private loan revenue is increasing on the feasible domain because $A - 2\gamma\bar{l} = 1/20 > 0$, and $t < a$ ensures that preparation beyond the recognition requirement reduces private profit; at the illustrated choice its marginal profit is $-349/13500$.

B Replication specification and binary readiness extension

All model and certificate equations use quantities normalised by reference assets; a reported bp gap is divided by 10,000 before applying (30). For bank reference assets M_i and normalised

Primitive and role	Value	Construction rationale
R, M: capacity and assessment	1, 2/3	Normalise reserved capacity and make strict recognition require $u = 2/3$.
D: payment demand	4/5	Leave positive delayed principal even after preparing the required amount.
$\bar{l}, \kappa$: loan cap and capital ratio	3/2, 1/5	Make one unit of preparation expenditure contract lending by five units, while preserving a positive loan book.
h: initial headroom	0	Select the binding-capital branch; inherited equity is $E = 3/10$.
γ: demand curvature	1/10	Retain diminishing marginal output and increasing private revenue throughout the loan domain.
A: demand intercept	7/20	Set marginal borrower output at the loan cap to $p_0 = 1/5 = \kappa$.
a, β: preparation cost	1/50, 1/200	Give positive marginal workload and $P(2/3) = 13/900$.
t: private timing compensation	1/200	Keep private delay savings below marginal preparation cost, so recognition determines readiness.
Δ: total payment loss	3/100	Lie strictly between the lower and upper limits in (27).

Table 4: A rational construction of the financing mechanism. Values select an analytically transparent economy satisfying the stated inequalities.

activity $z = Z/M_i$, the monetary resource law is $M_i C_z(Z/M_i) = \alpha_z Z + \beta_z Z^2/(2M_i)$. Thus common normalised curvature corresponds to monetary curvature inversely proportional to bank scale.

The positive local curvature coefficients are $h_l = .08$, $h_b = .035$, $h_k = .03$, $h_I = .30$, $h_q = .025$ and $h_d = .35$. Market slopes are $g_l = .065$ and $g_b = .018$. The borrowing fee is .015, compensated linear timing cost .09, downstream timing cost .02 and baseline linear issuance cost .05. Collateral advance factors span 65–90%, corresponding to haircuts of 35–10%. Preparation coefficients are the printed cost a multiplied by a class grid from .85 to 1.15, with $\beta_c = .12$. The finite strategy domain is

$$\begin{aligned} 0 \le l_i \le 2 \quad & \text{(loan opportunity bound)}, \\ 0 \le b_i \le 1.5 \quad & \text{(runnable-funding domain)}, \\ 0 \le I_i \le .6 \quad & \text{(equity-issuance domain)}, \\ 0 \le k_{ic}, u_{ic}, q_{isc} \le 1/C \quad & \text{(collateral-class domain)}, \\ 0 \le v_i \le .3 \quad & \text{(preparation budget)}, \\ 0 \le d_{is} \le 2 \quad & \text{(deferred-principal domain)}. \end{aligned}$$

The additional contractual and physical bounds in (3)–(7) also apply.

The terminal envelope uses the worst admissible aggregate loan and funding quantities, two and 1.5. The conservative loss coefficient on loans is $\max\{0, 2g_l + h_l - a_i\}$; the funding coefficient is $.008 + 1.5g_b + .75h_b$. Write these two coefficients as $\bar{c}_{li}$ and $\bar{c}_b$. The implemented terminal condition is

$$\begin{aligned} e_i \ge \bar{c}_{li} l_i + \bar{c}_b b_i + .003 f_i + o_s A_i \\ + (r_q + h_q/(2C)) \sum_c q_{isc} + (.09 + .35) d_{is} \quad \text{(terminal solvency)}. \end{aligned} \tag{33}$$

Positive securities surplus is omitted from this lower envelope. The row guarantees nonnegative terminal equity throughout the private strategy box. Unit tests exercise the adverse endpoints that determine its coefficients.

The producer exports a compressed numerical fixture containing the complete box, affine matrix, convex objective diagonal, two aggregate rank factors, epigraph coefficients, primal allocation and dual multipliers. An independently implemented shared verifier uses no optimiser. Exact input arithmetic operates on the stored binary rational coefficients that define each computed economy. Each result records source hashes, solver status, iterations, solve time, verifier time, the actual gap, conditioning, welfare intervals, active constraints and accounting reconciliation. The separate analytic verifier uses rational primitives directly.

The study programs generate the numerical fixtures, exact verification receipts, and figure and table inputs. The bank-local decomposition uses a batched separable convex problem and retains its full iterate history. A verifier implemented independently of the optimiser checks direct shareholder cashflows, complete accounts, spending projection, corrupted-certificate rejection, total reference mass, the rational financing witness, the terminal-solvency envelope, zero preparation cost and separation of facility capacity from HQLA.

A binary readiness extension asks whether the financing reversal also arises with fixed setup costs and private implementation. Its contract menu supports discrete preparation and hidden-action choices; the continuous model studies the marginal investment decision conditional on access. The binary exercise covers 240 point economies and 4,800 exact optimisation certificates. Independent checks also reconcile 34,560 dated physical ledgers and 960 resource accounts. The continuous headroom result in Proposition 3 follows from the assumptions and proof given in this paper.

C Complementary liquidity disclosures

Five matched Pillar 3 reports provide liquidity context for the June 2025 accounting data in Section 2. Their LCR measures are twelve-month averages ending in June, while NSFR is measured at that reporting date. Across these matched groups, LCR ranges from 130.4% to 156.0% and NSFR from 109.5% to 126.0%. These ratios describe the disclosed liquidity and stable-funding positions; the readiness model separately represents the preparation needed to turn collateral capacity into deadline cash.

Group	LCR (%)	NSFR (%)
Intesa Sanpaolo	145.0	121.4
Nordea	156.0	123.4
Santander	152.0	126.0
Deutsche Bank	132.7	120.4
BNP Paribas	130.4	109.5

Table 5: Matched June 2025 Pillar 3 disclosures. LCR averages and NSFR point measurements retain their respective reporting conventions.

Statements and Declarations

Data and code availability. The study code is confined to the model, numerical experiments, independent certificate checks, and construction of the reported figures and tables. The empirical inputs are public EBA and Pillar 3 disclosures, retained with normalized records and provenance. The code, data, numerical fixtures and verification receipts are available from the author.

Funding and competing interests. The author received no specific grant for this work. The author declares no competing interests.

Use of assistance. The ideas, the initial writing and the numerical experiments reported in this work are the author's own. Large language models were used only to polish the language of the manuscript. The author reviewed and edited all resulting text and takes full responsibility for the content of the publication. Scientific interpretation and approval remain the author's responsibility.

References

[1] Bank of England. Guide to pre-positioning loan collateral, 2026. URL https://www.bankofengland.co.uk/markets/eligible-collateral/guide-to-pre-positioning-loan-collateral. Updated July 15, 2026; accessed October 3, 2026.

[2] Basel Committee on Banking Supervision. Basel Framework, LCR30: High-quality liquid assets, 2020. URL https://www.bis.org/committees/bcbs/basel-framework/standard/lcr/30/inforce/2019-12-15/published/2020-06-05. Accessed September 9, 2026.

[3] Basel Committee on Banking Supervision. Basel Framework, LEV20: Leverage ratio calculation, 2020. URL https://www.bis.org/committees/bcbs/basel-framework/standard/lev/20/inforce/2023-01-01/published/2020-03-27. Published March 27, 2020; effective January 1, 2023; accessed October 3, 2026.

[4] Basel Committee on Banking Supervision. Basel Framework, RBC20: Calculation of minimum risk-based capital requirements, 2020. URL https://www.bis.org/committees/bcbs/basel-framework/standard/rbc/20/inforce/2023-01-01/published/2020-11-26. Published November 26, 2020; effective January 1, 2023; accessed October 3, 2026.

[5] Basel Committee on Banking Supervision. Basel Framework, NSF30: Available and required stable funding, 2024. URL https://www.bis.org/committees/bcbs/basel-framework/standard/nsf/30/inforce/2019-12-15/published/2024-07-05. Accessed September 9, 2026.

[6] J. Frédéric Bonnans and Alexander Shapiro. *Perturbation Analysis of Optimization Problems*. Springer, 2000. doi: 10.1007/978-1-4612-1394-9.

[7] Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004. doi: 10.1017/CBO9780511804441. URL https://web.stanford.edu/~boyd/cvxbook/.

[8] Giuseppe C. Calafiore. Budgeted robust intervention design for financial networks with common asset exposures. Technical Report 2603.27274, arXiv, 2026. URL https://arxiv.org/abs/2603.27274.

[9] Rodrigo Coelho, Mathias Drehmann, Diarmuid Murphy, and Ruth Walters. Navigating liquidity stress: Operational readiness for central bank support. FSI Insights 64, Bank for International Settlements, 2024. URL https://www.bis.org/publications/fsi-insight-64-navigating-liquidity-stress-operational-readiness-central-bank-support.

[10] Dean Corbae and Pablo D'Erasmo. Capital buffers in a quantitative model of banking industry dynamics. *Econometrica*, 89(6):2975–3023, 2021. doi: 10.3982/ECTA16930. URL https://www.philadelphiafed.org/-/media/FRBP/Assets/working-papers/2021/wp21-24.pdf.

[11] European Banking Authority. 2025 eu-wide transparency exercise: Full database and metadata, 2025. URL https://www.eba.europa.eu/risk-and-data-analysis/risk-analysis/eu-wide-transparency-exercise. June 2025 reference observations; accessed October 3, 2026.

[12] Kenneth A. Froot and Jeremy C. Stein. Risk management, capital budgeting, and capital structure policy for financial institutions: An integrated approach. *Journal of Financial Economics*, 47(1):55–82, 1998. doi: 10.1016/S0304-405X(97)00037-8. URL https://www.nber.org/papers/w5403.

[13] Grzegorz Hałaj and Sofia Priazhkina. Stressed but not helpless: Strategic behaviour of banks under adverse market conditions. Staff Working Paper 2021-35, Bank of Canada, 2021. URL https://www.bankofcanada.ca/2021/07/staff-working-paper-2021-35/.

[14] Dov Monderer and Lloyd S. Shapley. Potential games. *Games and Economic Behavior*, 14(1): 124–143, 1996. doi: 10.1006/game.1996.0044. URL https://doi.org/10.1006/game.1996.0044.

[15] Adriano A. Rampini and S. Viswanathan. Collateral, risk management, and the distribution of debt capacity. *The Journal of Finance*, 65(6):2293–2322, 2010. doi: 10.1111/j.1540-6261.2010.01616.x. URL https://scholars.duke.edu/publication/794624.

[16] Adriano A. Rampini, Amir Sufi, and S. Viswanathan. Dynamic risk management. *Journal of Financial Economics*, 111(2):271–296, 2014. doi: 10.1016/j.jfineco.2013.10.003. URL https://scholars.duke.edu/publication/1006436.

[17] Suresh Sundaresan and Kairong Xiao. Liquidity regulation and banks: Theory and evidence. *Journal of Financial Economics*, 151:103747, 2024. doi: 10.1016/j.jfineco.2023.103747. URL https://business.columbia.edu/sites/default/files-efs/citation_file_upload/Liquidity%20Regulation%20and%20Banks.pdf.