

Accounting Opacity, Zombie Externalities, and the Mandate for Fair-Value Disclosure

[Author names omitted for review]

April 2026

Abstract

When banks choose between available-for-sale (AFS, mark-to-market) and held-to-maturity (HTM, historical cost) accounting, a public-bad externality generates market failure: each solvent-distressed bank individually prefers HTM to avoid capital compliance costs, but their collective choice sustains insolvent zombie banks in the opaque pool at social cost Z that no individual bank internalizes. The mechanism operates through two competing forces. The zombie cost channel (δZ , discounted present value of misallocation) pulls toward a mandate by quantifying the social harm of opacity. The fiscal timing channel ($\Phi = (1 - \delta)[(1 - \rho) + \kappa_{\text{res}}]$) pulls against a mandate by measuring the cost of front-loading deposit insurance payouts that HTM defers. When the fiscal timing channel dominates ($\delta \leq \delta^{**} = K/(Z + K)$), mandatory disclosure is welfare-reducing for any level of insolvency. When the zombie cost channel dominates, the mandate is welfare-improving if and only if the pool insolvency fraction exceeds $p_{\text{ext}}^* = C/(\delta Z - \Phi + C)$. Higher resolution costs raise the mandate threshold rather than lower it, because they amplify the fiscal timing channel. Calibration against U.S. call report data (5,116 banks, 2020–2024) places aggregate insolvency at 4–9% during the 2022–2023 stress episode, consistently below $p_{\text{ext}}^* \approx 16.2\%$, consistent with the absence of an AFS mandate.

JEL Codes: G21, G28, M41, G12

Keywords: bank accounting, held-to-maturity, fair value, zombie banks, deposit insurance, welfare threshold, banking regulation

1 Introduction

Between 2020 and 2022, U.S. banks shifted roughly \$1 trillion in securities from available-for-sale (AFS) to held-to-maturity (HTM) accounting. Under AFS, securities are marked to market and unrealized losses flow through regulatory capital. Under HTM, securities are carried at historical cost and unrealized losses are invisible to depositors. As the Federal Reserve raised rates sharply in 2022, this accounting choice mattered: Silicon Valley Bank had reclassified \$91 billion into HTM before its collapse, obscuring from depositors the depth of its exposure. SVB’s supervisory data showed insolvency months before the run; its public financial statements did not. The bank failed not because information was unavailable, but because the accounting regime concealed it. The question this paper addresses is: why does accounting opacity persist as a market outcome when transparency is socially valuable, and when does the resulting market failure justify mandatory fair-value disclosure?

Accounting opacity in bank pools is a public-bad externality driven by two competing forces, discounted zombie cost and fiscal timing. The mandate threshold $p_{\text{ext}}^* = C/(\delta Z - \Phi + C)$ identifies when zombie cost dominates and the mandate is welfare-improving.

The mechanism. Three bank types coexist: insolvent zombies (Type I, fraction p), solvent-distressed banks with capital below the regulatory minimum (Type S, fraction Δ), and genuinely healthy banks (Type H). Healthy banks always choose AFS: disclosure of health costs nothing and HTM exposes them to run contagion when the opaque pool is sufficiently contaminated. Insolvent banks always choose HTM: AFS means immediate regulatory closure at zero payoff. The key decision belongs to solvent-distressed banks. Each Type S bank, taking others’ choices as given, strictly prefers HTM to AFS when the HTM pool is safe enough not to trigger a run. The individual payoff comparison is straightforward: F (franchise value, no compliance cost) under HTM versus $F - C$ (franchise value minus capital cost) under AFS. HTM dominates at the individual level.

This individual incentive generates a public-bad externality when the market plays the payoff-dominant equilibrium. When all Type S banks choose HTM, they dilute the pool’s insolvency fraction and sustain the no-run Nash equilibrium for depositors. Insolvent zombie banks survive in this pool. Each surviving zombie imposes a social cost Z on the broader economy through credit misallocation and deposit-spread erosion (Caballero et al., 2008; Acharya et al., 2019). No individual Type S bank internalizes this cost: zombie social harm does not enter any bank’s profit function. The market equilibrium under payoff dominance is a Nash equilibrium in which every Type S bank chooses HTM (Proposition 3). The market failure is a conditional claim: when the risk-dominance or basin-dominance criterion selects the AFS equilibrium instead, the market reaches the efficient outcome without intervention. Payoff dominance is the empirically relevant criterion when learning dynamics or industry

coordination favor the high-payoff outcome, as the 2020–2022 HTM reclassification wave suggests. The failure is a public-bad problem: individual banks have no incentive to deviate because the compliance cost C is borne privately while the zombie-elimination benefit $p\delta Z$ is shared.

Two competing forces. The extended welfare analysis reveals two forces that pull in opposite directions. The *zombie cost channel* (δZ , the discounted present value of zombie misallocation) pulls toward a mandate: the bigger the social harm of zombie persistence, the easier the mandate is to justify. The *fiscal timing channel* ($\Phi \equiv (1 - \delta)[(1 - \rho) + \kappa_{\text{res}}]$) pulls against a mandate: mandatory AFS disclosure front-loads deposit insurance and resolution costs that HTM defers, and the welfare planner must account for the cost of paying earlier. Both channels arise from the same model structure. Zombie costs are realized at date 4, so their present value at the mandate decision is $\delta Z < Z$. Resolution costs are paid at date 2 under AFS but not until date 4 under HTM, creating a timing wedge of size $\Phi > 0$.

The interaction between the two channels produces the paper’s central results. First, a critical discount factor $\delta^{**} = K/(Z + K)$ where $K = (1 - \rho) + \kappa_{\text{res}}$: when $\delta \leq \delta^{**}$, the fiscal timing channel dominates and mandatory disclosure is welfare-reducing for any level of insolvency. The standard narrative that accounting transparency should be tightened during banking stress has a necessary condition: the deposit insurance fund must be patient enough to absorb the front-loaded fiscal cost. Second, a mandate threshold $p_{\text{ext}}^* = C/(\delta Z - \Phi + C)$ above the naive benchmark $p^* = C/(Z + C)$: the two forces together require a larger zombie fraction before the mandate is justified than naive analysis suggests.

Resolution costs and the mandate threshold. A comparative static that runs against standard intuition follows from the fiscal timing channel. Higher resolution costs raise the mandate threshold rather than lower it: a larger κ_{res} amplifies Φ , making early AFS closure more costly and requiring a larger zombie fraction to offset the fiscal burden. The reversal holds when zombie costs Z and resolution costs κ_{res} are sufficiently independent. If entrenched zombie relationships drive both (so that $Z'(\kappa_{\text{res}}) > (1 - \delta)/\delta$), the sign can reverse. Remark 4.7 characterizes the condition under which the reversal holds, rather than asserting unconditional robustness.

At the baseline calibration for U.S. banking ($R = 1.05$, $\rho = 0.90$, $C = 0.01$, $Z = 0.06$, $\delta = 0.95$, $\kappa_{\text{res}} = 0.005$), the extended threshold is $p_{\text{ext}}^* \approx 16.2\%$ and the critical discount factor is $\delta^{**} \approx 0.636$. Call report data from 5,116 banks over 2020–2024 places the aggregate mark-to-market insolvency fraction at 4–9% throughout the 2022–2023 episode, well below p_{ext}^* . Section 6 presents this as a calibration and consistency check on the model’s mechanism, not as a causal test: the call report panel confirms that the model’s thresholds are in the

right ballpark for the observed episode and that the binary contamination test is better read as consistency evidence than as identification of the zombie externality channel.

Relation to the literature. Diamond and Dybvig (1983) establish the two-Nash-equilibria structure of the depositor coordination game for a single bank of known type. The present paper extends D-D in two directions: depositors face an opaque pool of heterogeneous banks, so the pool composition determines which equilibrium obtains; and banks strategically choose whether to enter the opaque pool, creating a market failure that D-D’s single-bank setup cannot produce.

Gao and Jiang (2018) model a single bank’s reporting discretion in a Diamond-Dybvig framework. Their planner designs disclosure for one bank; there is no accounting choice game among banks, no zombie externality, and no fiscal timing channel. The welfare threshold in Gao and Jiang depends on the probability that the bank is distressed and the cost of forced disclosure, structurally similar to $C/(Z + C)$, but arising from a single-bank planning problem rather than a multi-bank coordination failure. This paper’s extended threshold p_{ext}^* incorporates two additional economic forces (discounted zombie costs δZ and the fiscal timing parameter Φ) that Gao and Jiang’s single-bank setup cannot generate. The cross-bank contamination game is the source of both.

Plantin et al. (2008) show that fair-value accounting can destabilize banks by triggering fire-sale feedback loops in illiquid markets. This paper’s mechanism runs in a different direction: HTM opacity causes zombie persistence through pool contamination, not through mark-to-market feedback. Neither mechanism dominates unconditionally; the optimal accounting standard depends on which fragility mode applies.

Chen et al. (2024) study a single bank’s run risk when unrealized losses are large. The cross-bank contamination game (multiple banks choosing accounting regimes, with pool composition endogenous to those choices) and the welfare analysis with fiscal timing are absent from their framework. Jiang et al. (2024) and Granja et al. (2024) document the HTM reclassification wave empirically and show weaker banks reclassified more aggressively. This paper provides a formal equilibrium foundation for why that outcome arises and when it generates a welfare case for regulatory intervention.

Goldstein and Sapra (2013) and Bouvard et al. (2015) study voluntary versus mandatory transparency in bank-run environments. Goldstein and Sapra analyze the costs and benefits of stress test disclosure; Bouvard, Chaigneau, and de Motta show transparency is socially excessive in normal times but insufficient in crises. The present paper formalizes a parallel result through the mandate threshold p_{ext}^* : the market provides too little transparency precisely when crisis conditions make it most valuable. The difference in mechanism is that this paper’s market failure arises from an endogenous accounting choice game (with pool

contamination as the externality), not from exogenous information disclosure.

Boyarchenko et al. (2025) survey the gap between bank regulation theory and accounting practice, calling explicitly for unified formal models in which accounting choice is endogenous and welfare analysis incorporates fiscal and regulatory interactions. The model in Section 2 answers that call. The welfare threshold p_{ext}^* , the fiscal timing decomposition, and the critical discount factor δ^{**} address the gap Boyarchenko, Hachem, and Kleymenova identify.

Roadmap. Section 2 sets up the model: environment, timing, and agents’ problems. Section 3 derives the depositor equilibria and the Nash equilibrium of the accounting choice game. Section 4 presents the extended welfare analysis and the mandate thresholds, foregrounding the two-channel decomposition. Section 5 discusses implications, testable predictions, and limitations. Section 6 presents calibration and consistency checks against U.S. call report data. Section 7 concludes.

2 Model

2.1 Environment

A unit mass of banks indexed $b \in [0, 1]$ each holds a portfolio of long-duration assets funded by a continuum of depositors. Each bank is one of three types, drawn independently at date 0 from a commonly known distribution and observed only by the bank itself:

- **Type I (insolvent):** fraction $p \in (0, 1)$. Asset value is below total deposits at market prices. Type I fails with certainty at date 4 regardless of depositor actions (solvency failure).
- **Type S (solvent-distressed):** fraction $\Delta \in (0, 1 - p)$. Assets exceed deposits at market value, but mark-to-market equity is below the regulatory capital minimum. Type S survives if and only if the fraction of depositors withdrawing early satisfies $m < \kappa \in (0, 1)$.
- **Type H (healthy):** fraction $1 - p - \Delta \in (0, 1)$. Fully solvent and liquid; never fails regardless of depositor actions.

Deposit payoffs. Each depositor chooses between early withdrawal ($a = 1$) and late withdrawal ($a = 0$). Early withdrawal pays face value 1 with certainty. Late withdrawal pays $R > 1$ if the bank survives and $\rho \in [0, 1)$ if it fails.

Bank payoffs. Under **AFS** reporting, each bank’s type is publicly disclosed at date 2. Regulators close Type I banks at date 2 before depositors act; deposit insurance pays each depositor face value 1 and recovers ρ from the bank’s residual assets. Type S banks are disclosed as solvent and survive the depositor game, but must raise regulatory capital at shareholder cost $C > 0$. Type H banks are disclosed as healthy at no cost.

Under **HTM** reporting, types are not disclosed. All HTM-reporting banks are pooled from depositors’ perspective; a depositor at any HTM bank knows only the equilibrium distribution of types in the pool.

A bank that survives the depositor game without a full run earns franchise value $F > 0$ (deposit relationship rents from payment services, cross-selling, and below-market deposit pricing). A bank subject to a full run ($m = 1$) loses its depositor base entirely and earns franchise value 0, even if technically solvent.

Zombie costs. Type I banks that survive under HTM operate as zombie banks. Each surviving zombie imposes social cost $Z > 0$ per unit of assets on the broader economy. The two channels, following Caballero et al. (2008) and Acharya et al. (2019), are: (i) *credit misallocation*: zombie banks extend credit to connected non-viable firms at below-market rates to avoid recognizing losses, crowding out productive borrowers; and (ii) *deposit franchise erosion*: zombie banks offer above-market deposit rates to maintain funding, compressing the deposit spreads of healthy and solvent-distressed competitors. Both channels are sustained by the deposit insurance backstop that prevents depositor discipline. Because the externality flows through the deposit insurance subsidy, Z does not appear in any individual bank’s profit function. The paper treats Z as a reduced-form parameter; Section 5.3 discusses this limitation. Zombie costs are incurred during dates 1–4 and recognized at date 4, so their present value at date 2 is δZ , where $\delta \in (0, 1]$ is the discount factor between dates 2 and 4.

Fiscal costs. Normalizing each bank’s deposit base to one unit:

- Under **AFS** (Type I closed at date 2): the deposit insurance fund pays face value 1 to each depositor and recovers ρ from assets. Net DI payout per Type I bank: $(1 - \rho)$, incurred at date 2. Resolution administrative cost: $\kappa_{\text{res}} > 0$, also at date 2.
- Under **HTM** (zombie resolved at date 4): the DI fund tops up depositors who waited to face value 1, paying $(1 - \rho)$ per Type I bank. Resolution cost κ_{res} is also incurred at date 4.

The gross DI payout per Type I bank is $(1 - \rho)$ under both regimes. The mandate does not reduce total DI costs; it changes when they are paid. Define the *fiscal convenience*

parameter:

$$\Phi \equiv (1 - \delta)[(1 - \rho) + \kappa_{\text{res}}] \geq 0. \quad (1)$$

Φ measures the present-value fiscal saving to the DI fund from delaying Type I resolution. When $\delta = 1$ (no discounting), $\Phi = 0$.

Franchise rents and social welfare (Assumption A9). The franchise value F earned by a Type I zombie bank during its period of operation is a rent extracted by bank insiders (via compensation, dividends, or risk-shifting) from the bank’s future residual value. The accounting identity is as follows. The DI fund pays $(1 - \rho)$ per unit of deposits at Type I resolution, covering the gap between face value and asset recovery. This DI subsidy is what enables Type I banks to continue operating without depositor discipline: the zombie’s franchise rents flow to insiders precisely because depositors are insured and do not run. The private benefit (zombie insider rents = F per surviving zombie bank) is therefore financed by the public subsidy (DI payout = $(1 - \rho)$ per zombie), which the welfare formula already records as a cost. No separate social benefit from zombie franchise value enters welfare. Zombie franchise rents net to zero in social welfare: the $(1 - \rho)$ DI coverage accounts for the transfer, and Z captures the additional real-sector distortions (credit misallocation and deposit spread erosion) that lie beyond the private transfer.

What F includes. F is the net franchise value to a surviving bank: it includes the present value of deposit relationship rents from payment services, cross-selling, and below-market deposit pricing. It does *not* include depositor payoffs. Depositor payoffs (face value 1 on early withdrawal, or R on successful late withdrawal) are separate from the bank’s franchise value and enter welfare directly through the deposit payoff terms in the welfare formulas. The welfare expressions W_{ext}^* and $W_{\text{ext}}^{\text{AFS}}$ (Proposition 4) aggregate bank franchise values net of compliance costs and DI fiscal costs; depositor payoffs cancel across the planner’s calculation because all deposits are either repaid at face value or covered by deposit insurance.

2.2 Timing

The game has five dates:

Date 0 (type realization): Nature draws types independently from $(p, \Delta, 1 - p - \Delta)$. Each bank observes its own type privately.

Date 1 (accounting choice): Each bank independently and simultaneously chooses $r_b \in \{\text{AFS}, \text{HTM}\}$. Each bank knows its own type but not others’ types or choices.

Date 2 (disclosure): Accounting choices are publicly observed. AFS banks have their types disclosed; Type I banks are closed by the regulator. HTM banks are pooled: depositors

observe only that their bank reported HTM.

Date 3 (depositor game): All depositors simultaneously choose $a \in \{0, 1\}$. Depositors at AFS-disclosed banks know the bank's exact type. Depositors at HTM banks share the equilibrium posterior over HTM pool composition.

Date 4 (outcomes): Solvency failures are realized. Type I banks fail. Type S banks fail if and only if $m \geq \kappa$. Type H banks never fail. DI resolution costs, zombie costs, and AFS capital costs are paid.

2.3 Agents' Problems

Depositors. A depositor at bank b chooses $a_b \in \{0, 1\}$ to maximize expected payoff:

$$u(a, m, \tau_b) = a \cdot 1 + (1 - a)[\mathbf{1}\{\text{bank survives}\} \cdot R + \mathbf{1}\{\text{bank fails}\} \cdot \rho]. \quad (2)$$

At HTM banks, depositors form beliefs over pool composition. Let q denote the fraction of HTM banks that are Type I (pool insolvency fraction). Expected payoff to waiting given all others wait: $q\rho + (1 - q)R$.

Notation. The symbol δ denotes the discount factor between dates 2 and 4 (Assumption A10). In Lemma 1, the proof uses δ_S to denote the fraction of Type S banks within the HTM pool, distinct from the discount factor δ . The subscript S resolves the ambiguity within the lemma proof. All other uses of δ in the paper refer to the discount factor.

Banks. Bank b of type $\tau \in \{I, S, H\}$ chooses $r_b \in \{\text{AFS}, \text{HTM}\}$ at date 1 to maximize expected payoff, taking as given: (i) the equilibrium accounting strategies of all other banks, summarized by $\lambda \equiv$ fraction of Type S banks choosing HTM; (ii) the depositor equilibrium that follows from the resulting pool composition. Since each bank has measure zero, its own choice does not affect $q(\lambda) = p/(p + \lambda\Delta)$.

Regulator. The regulator optionally mandates AFS for all banks, maximizing social welfare $W_{\text{ext}}(\cdot)$ as defined in Section 4.

Equilibrium concept. The equilibrium is a rational expectations equilibrium: depositors form correct beliefs about equilibrium accounting choices and play optimally in the depositor game; banks choose accounting regimes anticipating depositors' best response to the equilibrium pool composition. By backward induction, the game has a unique subgame perfect equilibrium up to the equilibrium selection criterion in the accounting game (payoff dominance; see Appendix B).

2.4 Assumptions

The analysis proceeds under the following structural conditions:

A1 (Payoff ordering): $R > 1 > \rho \geq 0$.

A2 (Type fractions): $p \in (0, 1)$ and $\Delta \in (0, 1 - p)$.

A3 (Franchise vs. compliance): $0 < C < F$.

A4 (Zombie cost): $Z > 0$.

A5 (Withdrawal capacity): $\kappa \in (0, 1)$.

A6 (Genericity): For any pool insolvency fraction q arising in equilibrium, $q \neq 1 - t$. This rules out knife-edge mixed Nash equilibria in the depositor game.

A7 (Baseline mandate viability): $C(1 - \rho) < Z(R - 1)$. Under A7, $p^* \equiv C/(Z + C) < 1 - t$ (Lemma 2), so the baseline mandate threshold lies strictly below the opacity trap.

A7' (Extended mandate viability): $C(1 - \rho) < (R - 1)(\delta Z - \Phi)$. Under A7', $p_{\text{ext}}^* \equiv C/(\delta Z - \Phi + C) < 1 - t$ (Lemma 3). A7' is strictly stronger than A7 when $\Phi > 0$.

A8 (Fiscal viability): $\delta Z > \Phi$. Under A8, $p_{\text{ext}}^* \in (0, 1)$, so the mandate is welfare-improving for sufficiently high p . A8 holds at calibration ($\delta Z = 0.057 > \Phi = 0.00525$).

A9 (Zombie franchise as private transfer): Zombie franchise rents net to zero in social welfare; the $(1 - \rho)$ DI coverage already accounts for the transfer (see above).

A10 (Discounting convention): Dates 2 and 3 are equivalent for discounting purposes. The discount factor δ applies between dates 2–3 and date 4.

Define the *run threshold parameter*:

$$t \equiv \frac{1 - \rho}{R - \rho} \in (0, 1), \quad 1 - t = \frac{R - 1}{R - \rho}. \quad (3)$$

Under A1, $t \in (0, 1)$. The threshold $1 - t$ determines when depositors at an HTM pool find waiting at least as profitable as running: the no-run Nash equilibrium exists if and only if the pool insolvency fraction $q \leq 1 - t$ (Lemma 1).

3 Equilibrium Analysis

3.1 Depositor Game in the HTM Pool

For any HTM pool with insolvency fraction $q \in [0, 1)$ and Type S fraction δ_S within the pool, the following lemma characterizes depositor Nash equilibria.

Lemma 1 (HTM Depositor Equilibria).

- (a) $\sigma = 0$ (no run) is a Nash equilibrium if and only if $q \leq 1 - t$.
- (b) $\sigma = 1$ (run) is a Nash equilibrium if and only if $q + \delta_S \geq 1 - t$.
- (c) When $q > 1 - t$, $\sigma = 1$ is the **unique** Nash equilibrium.
- (d) When $q < 1 - t$ (strict), $\sigma = 0$ strictly Pareto dominates $\sigma = 1$ for depositors.

Proof. (a). Under $\sigma = 0$, a depositor's payoff from waiting is $q\rho + (1 - q)R$. No profitable deviation from waiting requires $q\rho + (1 - q)R \geq 1$, equivalently $(1 - q)(R - \rho) \geq 1 - \rho$, equivalently $q \leq 1 - t$.

(b). Under $\sigma = 1$, a depositor's payoff from waiting is $q\rho + \delta_S\rho + (1 - q - \delta_S)R$: Type I and Type S both fail (the latter because $m = 1 \geq \kappa$), while Type H pays R . No profitable deviation from running requires $(1 - q - \delta_S)(R - \rho) \leq 1 - \rho$, equivalently $q + \delta_S \geq 1 - t$.

(c). When $q > 1 - t$, part (a) shows $\sigma = 0$ is not an equilibrium. Consider a mixed Nash equilibrium $\sigma^* \in (0, 1)$. Depositors must be indifferent between running and waiting.

If $\sigma^* < \kappa$ (Type S survives), indifference requires $q\rho + (1 - q)R = 1$, i.e., $q = 1 - t$, contradicting $q > 1 - t$.

If $\sigma^* \geq \kappa$ (Type S fails), indifference requires $(1 - q - \delta_S)(R - \rho) = 1 - \rho$, i.e., $q + \delta_S = 1 - t$. But $q > 1 - t$ and $\delta_S \geq 0$, so $q + \delta_S > 1 - t$: contradiction. By A6, $\sigma^* = \kappa$ is also excluded.

Hence no mixed equilibrium exists; $\sigma = 1$ is unique.

(d). Under $\sigma = 0$, each depositor's payoff is $q\rho + (1 - q)R > 1$ (from strict inequality $q < 1 - t$). Under $\sigma = 1$, each depositor gets exactly 1. So $\sigma = 0$ strictly Pareto dominates $\sigma = 1$. □ □

Lemma 1(c) is the *opacity trap*: when the pool insolvency fraction exceeds $1 - t$, depositors run on every HTM bank regardless of its true solvency. The threshold $1 - t = (R - 1)/(R - \rho)$ increases in R (higher return makes waiting more attractive) and in ρ (better deposit insurance recovery reduces run incentive).

3.2 Individual Accounting Choices

Given the depositor equilibria in Lemma 1, each bank type's accounting choice reduces to a comparison of payoffs under AFS and HTM.

Proposition 1 (Type I Always Chooses HTM). *In any equilibrium of the accounting game, each Type I bank weakly prefers HTM to AFS, with strict preference whenever the HTM pool does not run.*

Proof. Under AFS, Type I is closed at date 2. Payoff = 0.

Under HTM, two cases arise. If $q(\lambda) > 1 - t$, the unique run equilibrium obtains (Lemma 1(c)): Type I fails by solvency and its payoff is 0, equal to AFS. If $q(\lambda) \leq 1 - t$, the no-run equilibrium obtains (Lemma 1(d)): Type I survives as a zombie with payoff $F > 0$. HTM weakly dominates, with strict dominance in the no-run case. \square \square

Proposition 2 (Type H Always Chooses AFS). *In any equilibrium, each Type H bank weakly prefers AFS to HTM. All Type H banks choose AFS.*

Proof. Under AFS, Type H is disclosed as healthy. Depositors do not run; Type H earns franchise value F .

Under HTM, two cases arise. If $q(\lambda) \leq 1 - t$, the no-run equilibrium obtains: Type H earns F , the same as under AFS. By a tiebreaking assumption, Type H chooses AFS when indifferent. If $q(\lambda) > 1 - t$, the unique run equilibrium obtains: all depositors run, Type H loses its depositor base entirely ($m = 1$), and earns franchise value 0. Since $F > 0$, AFS strictly dominates in this case.

Combined with the tiebreaking assumption in the no-run case, Type H always chooses AFS. \square \square

Pool composition. By Propositions 1 and 2, in any equilibrium the HTM pool contains all Type I banks and fraction $\lambda \in [0, 1]$ of Type S banks, with no Type H banks. The pool insolvency fraction is endogenous:

$$q(\lambda) = \frac{p}{p + \lambda\Delta}, \quad (4)$$

which is strictly decreasing in λ : more Type S banks dilute the pool's insolvency fraction.

3.3 The Accounting Externality

The zombie regime. The *zombie regime* holds when $q(1) = p/(p + \Delta) < 1 - t$: even if all Type S banks choose HTM, the pool remains below the run threshold. In the zombie regime, there exists a unique $\bar{\lambda} \in (0, 1)$ defined by $q(\bar{\lambda}) = 1 - t$, with closed form $\bar{\lambda} = pt/[\Delta(1 - t)]$ (Appendix A).

Proposition 3 (The Accounting Externality). *Suppose the zombie regime holds: $p/(p + \Delta) < 1 - t$.*

(a) **Individual best response.** For each Type S bank:

$$r_S^*(\lambda) = \begin{cases} AFS & \text{if } \lambda < \bar{\lambda}, \\ HTM & \text{if } \lambda > \bar{\lambda}, \\ \text{indifferent} & \text{if } \lambda = \bar{\lambda}. \end{cases} \quad (5)$$

(b) **Nash equilibria.** There are exactly two pure-strategy Nash equilibria: $\lambda = 0$ (all Type S choose AFS) and $\lambda = 1$ (all Type S choose HTM). The equilibrium $\lambda = 1$ strictly Pareto dominates $\lambda = 0$ for all Type S banks. The payoff-dominant equilibrium is $\lambda^* = 1$.

(c) **The externality.** Under $\lambda^* = 1$, each Type S bank earns F . But their collective choice sustains the HTM pool, allowing all Type I banks to survive as zombies and imposing social cost $p(\delta Z)$ on society. No individual Type S bank internalizes $p(\delta Z)$: zombie harm enters the social welfare function but not any bank's private payoff.

(d) **Social efficiency.** The market equilibrium $\lambda^* = 1$ is Pareto-dominant for Type S banks (each earns F instead of $F - C$). Whether it is socially efficient depends on the balance between zombie cost and fiscal timing effects; this is analyzed in Section 4, where Proposition 5(b) establishes the if-and-only-if condition $p > p_{ext}^*$.

Proof. (a). Each Type S bank takes λ as given (measure zero, does not affect $q(\lambda)$).

Under AFS: payoff = $F - C$, independent of λ .

Under HTM: if $\lambda < \bar{\lambda}$, then $q(\lambda) > 1 - t$ and the unique run equilibrium obtains. Type S fails ($m = 1 \geq \kappa$); payoff = 0. If $\lambda > \bar{\lambda}$, then $q(\lambda) < 1 - t$ and the no-run equilibrium obtains (Lemma 1(d)); Type S survives with payoff = F .

Since $F > F - C > 0$ (A3): HTM is preferred when $\lambda > \bar{\lambda}$; AFS is preferred when $\lambda < \bar{\lambda}$.

(b). A symmetric pure Nash equilibrium requires $\lambda^* \in \text{BR}(\lambda^*)$.

At $\lambda^* = 0$: $q(0) = 1 > 1 - t$. Each Type S's best response is AFS (payoff $F - C > 0$); any unilateral deviation to HTM yields payoff $0 < F - C$. No deviation.

At $\lambda^* = 1$: $q(1) < 1 - t$ (zombie regime). Each Type S's best response is HTM (payoff F); any deviation to AFS yields $F - C < F$. No deviation.

No other pure equilibrium exists: for $\lambda^* \in (0, \bar{\lambda})$, $\text{BR}(\lambda^*) = \{\text{AFS}\}$, inconsistent with $\lambda^* > 0$; for $\lambda^* \in (\bar{\lambda}, 1)$, $\text{BR}(\lambda^*) = \{\text{HTM}\}$, inconsistent with $\lambda^* < 1$; mixed equilibrium at $\lambda = \bar{\lambda}$ is excluded by A6.

Payoff comparison: under $\lambda = 1$ every Type S earns F ; under $\lambda = 0$ every Type S earns $F - C$. Since $F > F - C$ (A3), $\lambda = 1$ strictly Pareto dominates $\lambda = 0$. Appendix B discusses payoff dominance as the equilibrium selection criterion.

(c). Under $\lambda^* = 1$ with $q(1) < 1 - t$, depositors do not run. Type I banks survive as zombies. Each Type S bank's payoff is F ; the term $p(\delta Z)$ does not appear in any individual

bank's optimization problem. The social welfare function (Section 4) includes $p(\delta Z)$ as a cost; private payoffs do not.

(d). The Pareto dominance for Type S banks is immediate: $F > F - C$ (A3). Whether the outcome is socially efficient depends on the balance between zombie cost and fiscal timing; this is established in Proposition 5(b) in Section 4. \square \square

Economic interpretation. The externality is not a coordination failure. Banks coordinate *successfully* on HTM. The market failure is a public-bad problem: the compliance cost C is borne privately by each Type S bank that chooses AFS, while the zombie-elimination benefit $p\delta Z$ accrues to society as a whole. No individual bank has an incentive to deviate to AFS and pay the compliance cost to eliminate zombies that others free-ride on. The market reaches the bad outcome not despite coordination, but through it.

3.4 The Three-Regime Structure

Before the extended welfare analysis, consider the baseline (undiscounted, no fiscal costs: $\delta = 1$, $\kappa_{\text{res}} = 0$, so $\Phi = 0$). Social welfare under the market equilibrium is $W^* = (1-p)F - pZ$ and under an AFS mandate is $W^{\text{AFS}} = (1-p)(F - C)$. The baseline welfare gain from the mandate is:

$$G(p) \equiv W^{\text{AFS}} - W^* = pZ - (1-p)C, \quad (6)$$

with baseline mandate threshold $p^* = C/(Z + C)$.

Lemma 2 ($p^* < 1 - t$). *Under A7, $p^* \equiv C/(Z + C) < 1 - t \equiv (R - 1)/(R - \rho)$.*

Proof. $C/(Z + C) < (R - 1)/(R - \rho) \iff C(R - \rho) < (R - 1)(Z + C) \iff C(1 - \rho) < Z(R - 1)$. This is A7. \square \square

Lemma 3 ($p_{\text{ext}}^* < 1 - t$). *Under A7', $p_{\text{ext}}^* \equiv C/(\delta Z - \Phi + C) < 1 - t$.*

Proof. $C/(\delta Z - \Phi + C) < (R - 1)/(R - \rho) \iff C(1 - \rho) < (R - 1)(\delta Z - \Phi)$. This is A7'. \square \square

Lemmas 2 and 3 confirm the three-regime structure is non-empty: the zombie-mandate region $(p_{\text{ext}}^*, 1 - t)$ has positive measure. The extended welfare analysis in Section 4 derives p_{ext}^* formally.

4 Extended Welfare Analysis

The welfare analysis builds directly on the mechanism: when the market plays the payoff-dominant equilibrium, Type S banks collectively choose HTM and sustain zombie banks

at social cost Z . The planner's problem is to weigh the zombie cost channel against the fiscal timing channel, accounting for discounting and the front-loading of fiscal costs under AFS. All welfare expressions are per unit of bank assets. The two-type case ($\Delta = 1 - p$, no Type H banks) delivers clean closed-form formulas; Section 4.6 derives the three-type extension and shows it produces a stronger result.

4.1 Extended Welfare Formulas

Proposition 4 (Extended Welfare Formulas). *In the two-type case ($\Delta = 1 - p$) under the payoff-dominant equilibrium $\lambda^* = 1$ in the zombie regime ($p < 1 - t$):*

(a) *Market equilibrium:*

$$W_{ext}^* = (1 - p)F - p\delta Z - p\delta[(1 - \rho) + \kappa_{res}]. \quad (7)$$

(b) *AFS mandate:*

$$W_{ext}^{AFS} = (1 - p)(F - C) - p[(1 - \rho) + \kappa_{res}]. \quad (8)$$

(c) *Extended welfare gain:*

$$G_{ext}(p) \equiv W_{ext}^{AFS} - W_{ext}^* = p(\delta Z - \Phi) - (1 - p)C, \quad (9)$$

where $\Phi \equiv (1 - \delta)[(1 - \rho) + \kappa_{res}]$.

(d) *Decomposition:*

$$G_{ext}(p) = \underbrace{p\delta Z}_{\substack{\text{zombie cost} \\ \text{eliminated (PV)}}} - \underbrace{(1 - p)C}_{\substack{\text{compliance cost} \\ \text{imposed}}} - \underbrace{p\Phi}_{\substack{\text{fiscal timing} \\ \text{channel}}}. \quad (10)$$

Proof. (a). Under $\lambda^* = 1$, $q(1) = p < 1 - t$ (zombie regime). The no-run equilibrium obtains. Type I banks survive until date 4.

Type S (mass $1 - p$): Survives. Earns franchise value F . No DI cost or resolution. Contribution: $(1 - p)F$.

Type I (mass p): Survives dates 1–4 as zombie. Social cost Z recognized at date 4; present value at date 2 is δZ . Fails at date 4: DI pays $(1 - \rho)$ at date 4 (present value $\delta(1 - \rho)$); resolution cost κ_{res} at date 4 (present value $\delta\kappa_{res}$). By A9, zombie franchise rents net to zero in welfare. Contribution to W_{ext}^* : $p[-\delta Z - \delta(1 - \rho) - \delta\kappa_{res}]$.

Combining: (7).

(b). Under the AFS mandate, Type I (mass p) is closed at date 2. DI pays $(1 - \rho)$ at date 2 (present value $1 - \rho$); resolution cost κ_{res} at date 2. No zombie cost. Type S (mass $1 - p$): disclosed as solvent, no run, earns $F - C$, no DI cost. Combining: (8).

(c).

$$\begin{aligned}
G_{\text{ext}}(p) &= W_{\text{ext}}^{\text{AFS}} - W_{\text{ext}}^* \\
&= [(1 - p)(F - C) - p[(1 - \rho) + \kappa_{\text{res}}]] - [(1 - p)F - p\delta Z - p\delta[(1 - \rho) + \kappa_{\text{res}}]] \\
&= -(1 - p)C - p[(1 - \rho) + \kappa_{\text{res}}](1 - \delta) + p\delta Z \\
&= p(\delta Z - \Phi) - (1 - p)C. \quad \checkmark
\end{aligned}$$

(d). Immediate from (9) and $\Phi \geq 0$. □ □

Remark (regime invariance of direct DI costs). The gross DI payout per Type I bank is $(1 - \rho)$ under both AFS and HTM. The mandate does not reduce total deposit insurance costs; it changes only when they are paid. G_{ext} captures exactly two effects: zombie cost eliminated ($p\delta Z$) minus fiscal timing cost of mandate ($p\Phi$), net of compliance costs $((1 - p)C)$.

Remark (Timing convention for F and DI costs). The welfare formulas mix terms at different dates and it is worth stating the convention explicitly. Franchise value F represents operating income earned during the period ending at date 3, before date-4 resolution. Under Assumption A10, dates 2 and 3 are equivalent for discounting purposes, so F is already at the same present-value baseline as all date-2 quantities; no further discount applies. DI payout costs and resolution costs, by contrast, accrue at date 4 (under HTM) or date 2 (under AFS), which is why the HTM terms carry the factor δ and the AFS terms do not. Specifically: in equation (7), $(1 - p)F$ is at the date-2 baseline (Assumption A10), while $-p\delta[(1 - \rho) + \kappa_{\text{res}}]$ discounts date-4 costs back to date 2. In equation (8), $(1 - p)(F - C)$ is again at the date-2 baseline, while $-p[(1 - \rho) + \kappa_{\text{res}}]$ is undiscounted because AFS closure occurs at date 2. Discounting F consistently at δ would add $-p(1 - \delta)F$ to W_{ext}^* and $-(1 - p)(1 - \delta)F$ to $W_{\text{ext}}^{\text{AFS}}$, changing both welfare levels but not the sign or qualitative form of $G_{\text{ext}}(p) = W_{\text{ext}}^{\text{AFS}} - W_{\text{ext}}^*$, because the F terms cancel in the difference: the franchise value is earned by the same $(1 - p)$ mass of non-zombie banks under both regimes, so any uniform rescaling of F drops out of G_{ext} . All qualitative results and comparative statics are therefore invariant to this modeling choice; only the absolute welfare levels change.

4.2 Extended Mandate Threshold

Proposition 5 (Extended Mandate Threshold). *Under A1–A8:*

(a) **Threshold formula:**

$$p_{ext}^* \equiv \frac{C}{\delta Z - \Phi + C} \in (0, 1). \quad (11)$$

(b) **If-and-only-if condition:** $G_{ext}(p) > 0$ if and only if $p > p_{ext}^*$.

(c) **Comparison to baseline:** $p_{ext}^* > p^* \equiv C/(Z+C)$ whenever $\delta < 1$ and $(1-\rho) + \kappa_{res} > 0$. Equality holds if and only if $\delta = 1$.

(d) **Monotonicity:** p_{ext}^* is strictly increasing in Φ : higher fiscal advantage of HTM raises the mandate threshold.

(e) **Non-empty mandate region:** Under A7, $p_{ext}^* < 1 - t$, so the zombie-mandate regime $(p_{ext}^*, 1 - t)$ is non-empty.

Proof. (a). Under A8, $\delta Z - \Phi > 0$. With $C > 0$, $p_{ext}^* = C/(\delta Z - \Phi + C) \in (0, 1)$. \checkmark

(b). $G_{ext}(p) > 0 \iff p(\delta Z - \Phi) > (1-p)C \iff p(\delta Z - \Phi + C) > C \iff p > p_{ext}^*$.

\checkmark

(c). $\delta Z - \Phi = \delta Z - (1-\delta)[(1-\rho) + \kappa_{res}]$. Compare to Z :

$$Z - (\delta Z - \Phi) = (1-\delta)Z + (1-\delta)[(1-\rho) + \kappa_{res}] = (1-\delta)[Z + (1-\rho) + \kappa_{res}] > 0$$

when $\delta < 1$, $Z > 0$ (A4), and $(1-\rho) + \kappa_{res} > 0$ (A1). Hence $\delta Z - \Phi < Z$, giving $\delta Z - \Phi + C < Z + C$, and thus $p_{ext}^* > p^*$. When $\delta = 1$: $\Phi = 0$ and $\delta Z = Z$, so $p_{ext}^* = C/(Z + C) = p^*$.

(d). $\partial p_{ext}^*/\partial \Phi = C/(\delta Z - \Phi + C)^2 > 0$. \checkmark

(e). Lemma 3. □ □

Economic content: two channels, one threshold. The threshold formula p_{ext}^* is the algebraic implementation of the two-channel mechanism. The zombie cost channel enters through δZ in the denominator: larger zombie misallocation lowers the threshold, making the mandate easier to justify. The fiscal timing channel enters through Φ in the denominator: a larger timing wedge raises the threshold, making the mandate harder to justify. The formula captures the balance between the two channels, but the economic content is the mechanism, not the formula. A planner who simply wrote down the welfare difference and solved $G_{ext}(p) = 0$ would derive the same expression; the model's contribution is establishing why the market ends up at $\lambda^* = 1$ (the public-bad game structure) and characterizing the conditions under which that equilibrium generates a welfare loss worth correcting.

4.3 Extended vs. Baseline Welfare Gain: Decomposition and Channels

Proposition 6 (Extended vs. Baseline Welfare Gain). (a) For all $p \in (0, 1)$ and $\delta < 1$ with $(1 - \rho) + \kappa_{res} > 0$:

$$G_{ext}(p) < G(p) \equiv pZ - (1 - p)C. \quad (12)$$

(b) The gap decomposes as:

$$G(p) - G_{ext}(p) = p[(1 - \delta)Z + \Phi], \quad (13)$$

where $(1 - \delta)Z$ is the loss from discounting zombie benefits and Φ is the fiscal timing cost of accelerating DI payouts under AFS.

(c) The exact formula for $p_{ext}^* - p^*$ is:

$$p_{ext}^* - p^* = \frac{C[(1 - \delta)Z + \Phi]}{(\delta Z - \Phi + C)(Z + C)} > 0, \quad (14)$$

increasing in $(1 - \delta)$ and in Φ , decreasing in Z .

(d) Under HTM, no individual Type I bank internalizes the fiscal timing cost: each earns payoff $F > 0$ regardless of Φ . The aggregate fiscal timing cost $p\Phi$ is a second welfare-relevant distortion, separate from the zombie misallocation externality $p(\delta Z)$.

Proof. (a). $G(p) - G_{ext}(p) = pZ - p(\delta Z - \Phi) = p[(1 - \delta)Z + \Phi] > 0$ when $\delta < 1$ and $\Phi > 0$.

✓

(b). Direct from the decomposition. ✓

(c).

$$p_{ext}^* - p^* = C \cdot \frac{(Z + C) - (\delta Z - \Phi + C)}{(\delta Z - \Phi + C)(Z + C)} = \frac{C[(1 - \delta)Z + \Phi]}{(\delta Z - \Phi + C)(Z + C)}.$$

Positivity: $C > 0$, $(1 - \delta)Z > 0$, $\Phi > 0$, denominators > 0 under A8. ✓

(d). Each Type I bank's payoff under HTM is $F > 0$, independent of Φ . The DI fund's fiscal problem is not part of the bank's optimization. Φ enters only W_{ext} , not bank payoff functions. □ □

Proposition 6(d) identifies a second layer of market failure. The zombie misallocation externality (Z , not internalized by Type S banks) is the first layer. The fiscal timing cost (Φ , not internalized by Type I banks) is the second. Both layers support the case for a mandate; both are incorporated in p_{ext}^* .

Remark (Nature of the fiscal timing cost). Φ is a cost to the welfare planner, not a strategic externality in the game-theoretic sense. There is no strategic interaction between individual

banks and the DI fund: the DI fund’s problem does not enter any bank’s payoff function. Type I banks do not internalize Φ simply because the DI fund’s fiscal problem is not part of their optimization. This is a feature of the deposit insurance setup. The term “fiscal timing cost” (rather than “fiscal timing externality”) is the more precise label. The welfare planner must account for Φ when evaluating the mandate; no individual bank does.

4.4 Regime Dominance and the Critical Discount Factor

Proposition 7 (Regime Dominance and Critical Discount Factor). *Define $Z_\delta \equiv \delta Z / [(1 - \rho) + \kappa_{res}]$.*

(a) **Z-dominant regime.** *When $Z_\delta > 1$, A8 holds automatically and $p_{ext}^* < 1 - t$ under A7. The zombie externality is the primary welfare force.*

(b) **Φ -dominant regime.** *When $Z_\delta < 1$ and $\delta < 1$, there exists a critical discount factor:*

$$\delta^{**} \equiv \frac{(1 - \rho) + \kappa_{res}}{Z + (1 - \rho) + \kappa_{res}} \in (0, 1), \quad (15)$$

*such that A8 fails (equivalently $G_{ext}(p) < 0$ for all $p \in (0, 1)$) if and only if $\delta \leq \delta^{**}$. When A8 fails, mandatory AFS disclosure is welfare-reducing for any level of insolvency.*

(c) **Boundary.** *The mandate is viable (there exists p with $G_{ext}(p) > 0$) if and only if $\delta > \delta^{**}$. The critical δ^{**} is decreasing in $Z / [(1 - \rho) + \kappa_{res}]$: higher zombie costs expand the region where the mandate is viable.*

(d) **Comparative statics.** *In the viable regime (A8 holds):*

<i>Parameter</i>	<i>Effect on p_{ext}^*</i>	<i>Economic content</i>
$C \uparrow$	+	<i>Higher compliance cost raises the mandate threshold</i>
$Z \uparrow$	–	<i>Higher zombie cost lowers the threshold</i>
$\delta \uparrow$	–	<i>Patient DI fund reduces timing cost; also raises δZ</i>
$\kappa_{res} \uparrow$	+ (for $\delta < 1$)	<i>Higher resolution cost raises fiscal timing advantage of HTM</i>
$\rho \uparrow$	– (for $\delta < 1$)	<i>Better recovery reduces DI timing cost</i>

Proof. (a). When $\delta Z > (1 - \rho) + \kappa_{res}$, we have $\delta Z > (1 - \delta)[(1 - \rho) + \kappa_{res}] = \Phi$ (since $1 - \delta \leq 1$). A8 holds. \checkmark

(b). Let $K = (1 - \rho) + \kappa_{res} > 0$. A8 states $\delta Z > (1 - \delta)K$. A8 fails iff $\delta Z \leq (1 - \delta)K \iff \delta(Z + K) \leq K \iff \delta \leq K / (Z + K) = \delta^{**}$.

When A8 fails: $\delta Z - \Phi \leq 0$, so $G_{\text{ext}}(p) = p(\delta Z - \Phi) - (1 - p)C \leq -(1 - p)C < 0$ for all $p \in (0, 1)$. ✓

$\delta^{**} \in (0, 1)$ since $K > 0$ and $Z > 0$.

(c). $G_{\text{ext}}(p) > 0$ for some p iff $\delta Z - \Phi > 0$ iff $\delta > \delta^{**}$. $\partial \delta^{**} / \partial (Z/K) = \partial [1 / (1 + Z/K)] / \partial (Z/K) = -1 / (1 + Z/K)^2 < 0$. ✓

(d). Write $D = \delta Z - \Phi > 0$ (A8). Then $p_{\text{ext}}^* = C / (D + C)$.

$\partial p^* / \partial C = D / (D + C)^2 > 0$. $\partial p^* / \partial Z = -C\delta / (D + C)^2 < 0$. $\partial p^* / \partial \delta$: $\partial D / \partial \delta = Z + [(1 - \rho) + \kappa_{\text{res}}] > 0$, so $\partial p^* / \partial \delta < 0$. $\partial p^* / \partial \kappa_{\text{res}}$: $\partial D / \partial \kappa_{\text{res}} = -(1 - \delta) \leq 0$, so $\partial p^* / \partial \kappa_{\text{res}} = C(1 - \delta) / (D + C)^2 \geq 0$ (strictly positive for $\delta < 1$). $\partial p^* / \partial \rho$: $\partial D / \partial \rho = (1 - \delta) \geq 0$, so $\partial p^* / \partial \rho = -C(1 - \delta) / (D + C)^2 \leq 0$ (strictly negative for $\delta < 1$). □ □

Three comparative statics. The resolution cost result ($\partial p_{\text{ext}}^* / \partial \kappa_{\text{res}} > 0$ for $\delta < 1$) runs against standard intuition. Higher administrative resolution costs make early AFS closure more expensive: the mandate front-loads a fiscal cost that HTM defers. As resolution costs rise, the fiscal timing channel grows stronger, requiring a larger zombie fraction before the zombie cost channel offsets it. This effect vanishes when $\delta = 1$: in the undiscounted case, when payments happen does not matter. The result holds when Z and κ_{res} are treated as independent. Remark 4.7 characterizes when the sign reverses: if zombie costs and resolution costs move together (because entangled zombie relationships drive both), the standard intuition that costly resolution favors early disclosure can be restored. The result holds as a reduced-form finding conditional on sufficient independence between the two parameters, not as an unconditionally robust claim.

Part (b) characterizes the environments in which the mandate is unconditionally welfare-reducing. In high-inflation or high-discount-rate environments, δ falls toward δ^{**} . As $\delta \rightarrow \delta^{**}$, $p_{\text{ext}}^* \rightarrow \infty$ (the mandate threshold grows without bound) and the mandate eventually becomes welfare-reducing for all p . The standard narrative that accounting transparency should be tightened during banking stress has a necessary condition: the deposit insurance fund must be patient enough to absorb the front-loaded fiscal cost of AFS closure.

At the U.S. calibration, $\delta^{**} \approx 0.636$, implying the mandate would be unconditionally welfare-reducing only at real discount rates above roughly 57% per year, well outside the relevant range. The mandate remains viable at U.S. parameter values.

4.5 Extended Three-Regime Structure

Proposition 8 (Extended Three-Regime Structure). *Under A1–A8 and A7' (two-type case), the equilibrium and welfare outcomes are fully characterized by the thresholds $p_{\text{ext}}^* < 1 - t$:*

<i>Regime</i>	<i>Condition</i>	<i>Market λ^*</i>	<i>Mandate verdict</i>
<i>Deep zombie</i>	$p < p_{ext}^*$	1	<i>Welfare-reducing ($G_{ext} < 0$)</i>
<i>Zombie-mandate</i>	$p_{ext}^* < p < 1 - t$	1	<i>Welfare-improving ($G_{ext} > 0$)</i>
<i>Opacity trap</i>	$p > 1 - t$	0	<i>Redundant ($G_{ext} = 0$)</i>

Proof. Market equilibrium λ^* follows from Propositions 1–3: fiscal parameters do not enter bank payoff functions. Welfare gains $G_{ext}(p)$ follow from Proposition 5(b) for $p \leq 1 - t$.

In the opacity trap ($p > 1 - t$): $q(\lambda) > 1 - t$ for any $\lambda < \bar{\lambda}$; the unique run equilibrium obtains regardless of accounting choice. Type S voluntarily chooses AFS (Proposition 2’s argument applies: HTM leads to a run, AFS does not). The AFS mandate reproduces the market outcome exactly: $G_{ext} = 0$. \square \square

Remark (regime structure). The extended three-regime structure narrows the zombie-mandate zone relative to the baseline. Standard welfare analyses ignoring fiscal timing and discounting of zombie costs use p^* rather than p_{ext}^* . Because $p_{ext}^* > p^*$, they overstate the insolvency range for which the mandate is justified.

Remark (Opacity trap welfare, including Type H). The regime table reports $G_{ext} = 0$ in the opacity trap ($p > 1 - t$). This claim requires explicit verification for all three bank types.

Market equilibrium in the opacity trap. When $p > 1 - t$, the pool insolvency fraction $q(\lambda) > 1 - t$ for any $\lambda < \bar{\lambda}$; the unique depositor equilibrium is a run on any HTM bank (Lemma 1(c)). Because HTM triggers a run, every solvent bank strictly prefers AFS. By Proposition 2, Type H banks always choose AFS (they are never capital-constrained and AFS weakly dominates for them). Type S banks also switch to AFS voluntarily: the run cost under HTM exceeds the compliance cost C . The market equilibrium is therefore $\lambda^* = 0$ with all banks in AFS.

Welfare enumeration under the market outcome ($\lambda^ = 0$).*

- **Type H banks** (mass $1 - p - \Delta$): choose AFS voluntarily; disclosed as healthy; no run; earn franchise value F . DI cost: 0.
- **Type S banks** (mass Δ): choose AFS voluntarily; disclosed as solvent; earn $F - C$. DI cost: 0.
- **Type I banks** (mass p): disclosed as insolvent under AFS; closed at date 2. DI pays $(1 - \rho)$ at date 2; resolution cost κ_{res} at date 2.

Market welfare: $W_{trap}^{mkt} = \Delta(F - C) + (1 - p - \Delta)F - p[(1 - \rho) + \kappa_{res}] = (1 - p)F - \Delta C - p[(1 - \rho) + \kappa_{res}]$.

Welfare enumeration under the AFS mandate.

- **Type H banks:** already in AFS; mandate is redundant for them. Earn F , same as under the market outcome.
- **Type S banks:** already in AFS under the market outcome; mandate does not change their accounting choice or payoff. Earn $F - C$.
- **Type I banks:** closed at date 2, same as under the market outcome. DI pays $(1 - \rho) + \kappa_{res}$ at date 2.

Mandate welfare: $W_{trap}^{AFS} = W_{trap}^{mkt}$.

Welfare gain: $G_{ext} = W_{trap}^{AFS} - W_{trap}^{mkt} = 0$.

The result is exact, not approximate. The mandate is redundant in the opacity trap because the market has already solved the problem: every bank type independently reaches the AFS outcome. Type H banks are the clearest case: Proposition 2 establishes they always choose AFS, so the mandate is irrelevant for them under both regimes. The welfare formula $G_{ext} = 0$ holds for each type individually and therefore for the aggregate.

The three-type analysis below confirms that the regime structure is preserved when Type H banks are present, with threshold $p_{ext}^{*,3} = \Delta C / (\delta Z - \Phi + \Delta C) < p_{ext}^*$. The two-type case is conservative: the mandate is justified at lower insolvency fractions in the full three-type model.

4.6 Three-Type Welfare Analysis

Proposition 9 (Three-Type Extended Welfare). *In the three-type case with $p, \Delta \in (0, 1)$ and $p + \Delta < 1$, under the payoff-dominant equilibrium $\lambda^* = 1$ in the zombie regime:*

(a) **Market equilibrium welfare:**

$$W_{ext}^{*,3} = (1 - p)F - p\delta Z - p\delta[(1 - \rho) + \kappa_{res}]. \quad (16)$$

(b) **AFS mandate welfare:**

$$W_{ext}^{AFS,3} = (1 - p)F - \Delta C - p[(1 - \rho) + \kappa_{res}]. \quad (17)$$

(c) **Three-type extended welfare gain:**

$$G_{ext}^3(p) = p(\delta Z - \Phi) - \Delta C. \quad (18)$$

(d) **Three-type mandate threshold:**

$$p_{ext}^{*,3} = \frac{\Delta C}{\delta Z - \Phi + \Delta C}. \quad (19)$$

$G_{\text{ext}}^3(p) > 0$ if and only if $p > p_{\text{ext}}^{*,3}$.

- (e) **Comparison to two-type threshold.** $p_{\text{ext}}^{*,3} < p_{\text{ext}}^*$ whenever $\Delta < 1$ (i.e., whenever Type H banks are present). The mandate is justified at a strictly lower insolvency fraction in the three-type model: the two-type welfare analysis is conservative.
- (f) **Regime structure preserved.** The three-regime structure of Proposition 8 holds with $p_{\text{ext}}^{*,3}$ replacing p_{ext}^* . Comparative statics on $p_{\text{ext}}^{*,3}$ retain the same signs as in Proposition 7(d). The zombie-mandate region $(p_{\text{ext}}^{*,3}, 1 - t)$ is non-empty under A7' (replacing C with ΔC).

Proof. (a). Under $\lambda^* = 1$, the no-run equilibrium obtains. Type I banks (mass p) survive as zombies until date 4, incurring social cost δZ each and generating DI costs $\delta[(1 - \rho) + \kappa_{\text{res}}]$ at date 4. Type S banks (mass Δ) earn franchise value F with no DI cost. Type H banks (mass $1 - p - \Delta$) earn F (they chose AFS voluntarily; no run; no compliance cost because they are not capital-constrained). Social welfare aggregates:

$$\begin{aligned} W_{\text{ext}}^{*,3} &= \Delta F + (1 - p - \Delta)F - p\delta Z - p\delta[(1 - \rho) + \kappa_{\text{res}}] \\ &= (1 - p)F - p\delta Z - p\delta[(1 - \rho) + \kappa_{\text{res}}]. \end{aligned}$$

(b). Under the AFS mandate, Type I (mass p) is closed at date 2: DI pays $(1 - \rho) + \kappa_{\text{res}}$ at date 2. Type S (mass Δ) is disclosed as solvent: earns $F - C$. Type H (mass $1 - p - \Delta$) is voluntarily AFS and earns F with no compliance cost.

$$\begin{aligned} W_{\text{ext}}^{\text{AFS},3} &= \Delta(F - C) + (1 - p - \Delta)F - p[(1 - \rho) + \kappa_{\text{res}}] \\ &= (1 - p)F - \Delta C - p[(1 - \rho) + \kappa_{\text{res}}]. \end{aligned}$$

(c). Subtract (16) from (17):

$$\begin{aligned} G_{\text{ext}}^3 &= -\Delta C - p[(1 - \rho) + \kappa_{\text{res}}] + p\delta Z + p\delta[(1 - \rho) + \kappa_{\text{res}}] \\ &= p\delta Z - p(1 - \delta)[(1 - \rho) + \kappa_{\text{res}}] - \Delta C \\ &= p(\delta Z - \Phi) - \Delta C. \quad \checkmark \end{aligned}$$

(d). $G_{\text{ext}}^3(p) > 0 \iff p(\delta Z - \Phi) > \Delta C \iff p > \Delta C / (\delta Z - \Phi + \Delta C) = p_{\text{ext}}^{*,3}$. Under A8, $\delta Z - \Phi > 0$ and $p_{\text{ext}}^{*,3} \in (0, \infty)$. For $p_{\text{ext}}^{*,3} < 1$, a sufficient condition is $\Delta C < \delta Z - \Phi$, equivalently $C(1 - \rho)\Delta < (R - 1)(\delta Z - \Phi)$.

(e). $p_{\text{ext}}^{*,3}/p_{\text{ext}}^* = \Delta C/C = \Delta < 1$ when $\Delta < 1$. \checkmark

(f). Write $D = \delta Z - \Phi > 0$. Then $p_{\text{ext}}^{*,3} = \Delta C / (D + \Delta C)$. All comparative statics follow identically to Proposition 7(d) with C replaced by ΔC : the signs are unchanged. The market

equilibrium and regime boundaries follow from Propositions 1–3, which do not depend on Δ . □ □

Economic content: the three-type case produces a stronger result. The compliance cost under the mandate is borne only by the Δ fraction of Type S banks, not by the entire $(1 - p)$ mass of non-zombie banks. Type H banks choose AFS voluntarily and pay no compliance cost. As a result, the threshold $p_{\text{ext}}^{*,3}$ is strictly lower than the two-type threshold p_{ext}^* : the mandate is justified at a smaller insolvency fraction when Type H banks are present, because fewer banks bear the compliance cost. Quantitatively, at baseline calibration with $\Delta = 0.85$: $p_{\text{ext}}^{*,3} \approx \Delta C / (\delta Z - \Phi + \Delta C) \approx 0.0085 / 0.06025 \approx 14.1\%$, compared to the two-type value of 16.2%. The zombie-mandate region expands modestly in the three-type case. All qualitative regime comparisons in the paper are conservative relative to the three-type case.

4.7 Calibration

The paper uses the following baseline parameters, consistent with the SVB-era U.S. banking system: $R = 1.05$, $\rho = 0.90$, $C = 0.01$, $Z = 0.06$, $F = 0.15$, $\delta = 0.95$, $\kappa_{\text{res}} = 0.005$.

Derived quantities:

$$\begin{aligned} t &= (1 - 0.90) / (1.05 - 0.90) = 2/3. \quad 1 - t = 1/3 \approx 0.333. \\ \delta Z &= 0.95 \times 0.06 = 0.057. \\ \Phi &= 0.05 \times (0.10 + 0.005) = 0.00525. \\ \delta Z - \Phi &= 0.05175. \quad (\text{A8: } 0.05175 > 0. \checkmark) \\ p_{\text{ext}}^* &= 0.01 / 0.06175 \approx 0.162. \quad p^* = 0.01 / 0.07 \approx 0.143. \\ \delta^{**} &= 0.105 / (0.06 + 0.105) \approx 0.636. \end{aligned}$$

Assumption A7' verification: $(R - 1)(\delta Z - \Phi) = 0.05 \times 0.05175 = 0.002588 > C(1 - \rho) = 0.001$. ✓

At $p = 0.20$ (illustrative counterfactual, above p_{ext}^* ; the 2022–2023 episode never reached this level):

$$G_{\text{ext}}(0.20) = 0.20 \times 0.05175 - 0.80 \times 0.01 = 0.01035 - 0.008 = 0.00235 \text{ per dollar of assets.}$$

Applied to \$6.9 trillion in HTM-intensive bank assets, the counterfactual welfare gain at $p = 0.20$ is approximately \$16.2 billion per year. This figure is purely illustrative: it quantifies the welfare gain the mandate would deliver if insolvency were at 20%, a threshold the 2022–2023 episode did not reach.

The baseline (undiscounted) counterfactual welfare gain at $p = 0.20$ is $G(0.20) = 0.004$, equivalent to approximately \$27.6 billion per year at the same counterfactual level. The fiscal timing correction reduces the estimated welfare gain by \$11.4 billion per year, decomposed as: discounting of zombie benefits ($p(1 - \delta)Z$ per dollar \approx \$4.1B) and fiscal timing cost ($p\Phi$ per dollar \approx \$7.2B). The two channels are quantitatively non-trivial and together reverse the welfare sign at insolvency fractions between $p^* \approx 14.3\%$ and $p_{\text{ext}}^* \approx 16.2\%$.

Quantitative role of the fiscal timing channel. At baseline calibration, $\Phi \approx 0.00525$ and $\delta Z \approx 0.057$, so $\Phi/(\delta Z) \approx 0.09$: the fiscal timing channel is approximately 9% the size of the zombie cost channel. This is a small but economically meaningful correction. The threshold shifts from $p^* \approx 14.3\%$ to $p_{\text{ext}}^* \approx 16.2\%$, a 13% increase in the mandate threshold that determines whether intervention is warranted. The channel matters not because it dominates the zombie cost in normal operating ranges, but because of two properties that a first-order welfare calculation would miss: its direction (it raises the mandate threshold, making it harder to justify intervention) and the sign reversal it produces in the κ_{res} comparative static. The regime in which fiscal timing unconditionally dominates ($\delta \leq \delta^{**} \approx 0.636$, implying annual real discount rates above roughly 57%) is, as stated, well outside the U.S. range. In economies with higher discount rates, larger resolution costs, or lower recovery rates, the quantitative weight of Φ grows. For instance, at $\delta = 0.85$ (consistent with moderate inflation or high sovereign risk in emerging markets), Φ rises to $(1 - 0.85)(0.10 + 0.005) \approx 0.0158$, roughly three times the U.S. baseline value, and $\Phi/(\delta Z) \approx 0.31$. The channel is not negligible in parameter regions relevant to banking systems outside the United States.

Remark (Z and κ_{res} correlation). The paper treats Z and κ_{res} as independent parameters. In practice, high resolution costs may signal more entrenched zombie relationships: difficult resolutions tend to involve banks with more distorted balance sheets, which in turn generate more zombie lending. If $Z = Z(\kappa_{\text{res}})$ with $Z'(\kappa_{\text{res}}) > 0$, the total derivative of p_{ext}^* with respect to κ_{res} becomes:

$$\frac{dp_{\text{ext}}^*}{d\kappa_{\text{res}}} = \frac{C(1 - \delta) - C\delta Z'(\kappa_{\text{res}})}{(\delta Z - \Phi + C)^2}. \quad (20)$$

The sign flips from positive to negative when $Z'(\kappa_{\text{res}}) > (1 - \delta)/\delta$. At baseline calibration ($\delta = 0.95$), this requires $Z'(\kappa_{\text{res}}) > 0.0526$ per unit of resolution cost. If a one-unit increase in resolution cost raises zombie losses by more than 5.26 cents on the dollar of bank assets, the standard intuition (costly resolution favors early disclosure) is restored. The comparative static on κ_{res} is therefore a reduced-form result that holds when Z and κ_{res} are sufficiently independent.

5 Discussion

5.1 Testable Predictions

The model generates several predictions that extend beyond the 2022–2023 calibration.

P1 (Aggregate insolvency threshold test). The model predicts no welfare-improving mandate when $p < p_{\text{ext}}^*$. Cross-country variation in deposit insurance coverage (ρ), resolution cost structures (κ_{res}), and discount rates (δ) generates cross-country variation in p_{ext}^* . Countries with higher resolution costs (lower ρ , higher κ_{res}) require a larger insolvency fraction before the mandate is justified. European banking systems, where zombie costs have been estimated at 5–10% by Acharya et al. (2019), would have p_{ext}^* approximately 3–4 percentage points lower than the U.S. estimate, making the mandate viable at lower observed stress levels.

P2 (Interest-rate channel). When interest rates rise, both the bank asset return R and the pool insolvency fraction p shift. The rise in p through mark-to-market losses pushes the system toward the zombie-mandate threshold, while the rise in R expands the opacity trap boundary $1 - t$. The net effect depends on whether the insolvency elasticity $p'(r)$ is large relative to the return channel $\partial(1 - t)/\partial R \cdot R'(r)$. The HTM surge of 9 percentage points (2020–2022) is qualitatively consistent with the first channel dominating at U.S. observed rates.

P3 (Fiscal discount environment). In high-inflation environments (low δ), the model predicts the mandate threshold rises sharply and the viable mandate region shrinks. Mandatory fair-value reporting is harder to justify during inflationary episodes (the opposite of the standard zombie narrative). This generates cross-time variation in optimal accounting policy that does not appear in models ignoring fiscal timing. Specifically, p_{ext}^* is strictly decreasing in δ (Proposition 7(d)), so periods with higher real discount rates require higher insolvency fractions to justify the mandate.

P4 (Resolution cost prediction). Jurisdictions with higher administrative bank resolution costs (κ_{res}) exhibit higher effective mandate thresholds. Cross-country variation in resolution cost estimates from FDIC annual reports and EU Single Resolution Board data provides empirical variation in p_{ext}^* . This direction (higher resolution costs raise the mandate threshold) is testable because it runs against standard intuition that costly resolution favors early disclosure.

5.2 Relationship to Existing Results

Gao and Jiang (2018). Gao and Jiang (2018) model a single bank’s reporting discretion in a Diamond-Dybvig framework. Their social planner designs disclosure for one bank facing a probability of distress, with a disclosure cost structurally similar to C and an opacity benefit from concealing bad news. The planner’s welfare condition produces a threshold on the insolvency probability above which disclosure is optimal. The threshold is analogous in form to $p^* = C/(Z + C)$.

The differences are structural. Gao and Jiang have a single bank; there is no accounting choice game among banks and no pool contamination mechanism. Without pool contamination, there is no zombie externality: in a single-bank model, the disclosure decision of one bank does not affect depositors at other banks. The extended threshold $p_{\text{ext}}^* = C/(\delta Z - \Phi + C)$ incorporates two economic forces that Gao and Jiang’s setup cannot generate: the discounting of zombie benefits (δZ , arising because zombies impose costs at date 4 but the mandate decision is at date 1) and the fiscal timing channel (Φ , arising because the multi-bank setup has a DI fund that bears the cost of earlier resolution under AFS). These are not relabelings of Gao and Jiang’s parameters; they arise from the cross-bank game structure.

Plantin, Sapra, and Shin (2008). Plantin et al. (2008) show that fair-value accounting destabilizes banks through fire-sale feedback: when marking to market forces asset sales, prices fall, forcing further sales. This paper’s mechanism is orthogonal. HTM opacity causes zombie persistence through pool contamination: insolvent banks survive because depositors cannot identify them, not because marking to market triggers liquidation. In PSS, the problem is that AFS is too transparent and causes destabilizing feedback; here the problem is that HTM is insufficiently transparent and sustains zombies. Neither mechanism dominates unconditionally; the optimal accounting standard depends on which fragility mode is more severe in context.

Chen, Yi, and Zhang (2024). Chen et al. (2024) study a single bank’s exposure to run risk when unrealized losses are large, driven by rate increases. They show that HTM accounting, by concealing unrealized losses, increases run risk for that single bank. The paper extends beyond Chen, Yi, and Zhang in two directions. First, the cross-bank contamination game: multiple banks choose accounting regimes, and the pool composition is endogenous to those choices. A single bank choosing AFS does not eliminate the zombie externality; all Type S banks must switch for the pool to be clean. Second, the welfare analysis with fiscal timing and the extended threshold p_{ext}^* are absent from their framework. Chen, Yi, and Zhang have no welfare threshold and no fiscal externality.

Boyarchenko, Hachem, and Kleymenova (2025). Boyarchenko et al. (2025) survey the gap between bank regulation theory and accounting practice. They identify the absence of formal models in which accounting choice is endogenous, depositor equilibria depend on the resulting pool composition, and welfare analysis incorporates fiscal and regulatory timing. The model in Section 2 provides such a framework. The welfare threshold p_{ext}^* , the two-channel decomposition (zombie misallocation cost and fiscal timing cost), and the critical discount factor δ^{**} address the gap Boyarchenko, Hachem, and Kleymenova identify.

Goldstein and Sapra (2013) and Bouvard, Chaigneau, and de Motta (2015). Goldstein and Sapra (2013) analyze the costs and benefits of disclosing bank stress test results in a Diamond-Dybvig-style environment. Their central trade-off (disclosure reduces depositor uncertainty but may trigger runs by revealing bad news) bears on the mandated versus voluntary AFS choice studied here. The key difference is that Goldstein and Sapra study disclosure of an exogenous regulator signal, while this paper studies banks' endogenous choice of accounting regime: the pool composition and the run threshold are jointly determined by the equilibrium accounting strategies. Bouvard et al. (2015) study voluntary versus mandatory transparency in a rollover-risk model and show that transparency is socially excessive in normal times but insufficient in crises. The present paper's mandate threshold p_{ext}^* formalizes a similar idea in an accounting-choice framework: the market produces too little transparency (HTM) precisely when crisis conditions ($p > p_{\text{ext}}^*$) make transparency most valuable. A technical difference is that Bouvard, Chaigneau, and de Motta study continuous investment rollover while this paper uses discrete Diamond-Dybvig runs, which generates the sharp threshold structure.

Heider, Hoerova, and Holthausen (2015). Heider et al. (2015) show that opacity in interbank markets leads to liquidity hoarding when banks cannot observe counterparty quality. The mechanism here operates through the depositor game rather than interbank credit, and generates zombie persistence rather than hoarding. The two models are complementary: opacity imposes costs at both the retail depositor and interbank levels.

5.3 Limitations

Equilibrium selection. The market failure result holds when the market plays the payoff-dominant equilibrium $\lambda^* = 1$. When risk dominance or basin dominance criteria select $\lambda^* = 0$ instead, the market naturally reaches the efficient outcome without regulatory intervention, and the welfare comparison in Section 4 does not apply. This is a scope condition on the paper's normative claims, not a flaw in the model: the paper characterizes when a mandate is welfare-improving conditional on HTM coordination obtaining.

Within the zombie-mandate regime, the basin-of-attraction argument (Appendix C) supports $\lambda^* = 1$ in the lower portion (where the basin condition $\bar{\lambda} < 1/2$ holds) but fails for $p \geq 0.20$ in the two-type case. Payoff dominance ($\lambda = 1$ Pareto dominates $\lambda = 0$ for all Type S banks) is the paper’s primary selection criterion, and it holds throughout the zombie-mandate regime regardless of basin size. The welfare results in Section 4 are conditional on payoff dominance.

Two observations support payoff dominance as the empirically relevant criterion. First, industry coordination mechanisms (association norms, peer bank behavior, regulatory lag) favor convergence to the high-payoff equilibrium: the 2020–2022 HTM reclassification wave occurred without any coordinating mandate, consistent with banks converging to HTM through learning dynamics. Second, in historical banking episodes with zombie lending (Japan 1990s, Europe 2012–2016), HTM-equivalent accounting persisted for years, suggesting the payoff-dominant equilibrium is self-reinforcing once established.

A full resolution via global games (Goldstein and Pauzner, 2005) would uniquely select the equilibrium and sharpen the comparative statics. The main welfare results do not require the global-games structure.

Exogenous zombie cost. The paper treats Z as a reduced-form parameter, following Caballero et al. (2008). Calibration uses zombie cost estimates from that literature (Japan: 3–6% of assets; Europe: 5–10%). Z is not derived from primitives in this model. A two-sector model in which zombie lending endogenously crowds out productive investment, along the lines of Acharya et al. (2019), would generate tighter predictions on how Z varies with observable bank characteristics (loan portfolio composition, borrower quality, market structure). That extension is not pursued; the welfare results here hold for any $Z > 0$ satisfying A4 and A7’.

Two-type versus three-type welfare analysis. The main welfare analysis (Propositions 4–8) uses the two-type case ($\Delta = 1 - p$, no Type H banks) for closed-form tractability. Proposition 9 (Section 4.6) derives the full three-type welfare formulas and proves the regime structure is preserved, with threshold $p_{\text{ext}}^{*,3} = \Delta C / (\delta Z - \Phi + \Delta C)$ strictly below the two-type threshold p_{ext}^* whenever Type H banks are present. The three-type case is in fact the stronger result: compliance costs fall only on the Δ fraction of Type S banks, making the mandate easier to justify. At baseline calibration, $p_{\text{ext}}^{*,3} \approx 14.1\%$, approximately 2 percentage points below the two-type value. The two-type results are conservative.

AFS closure assumption. The model assumes AFS disclosure leads to pre-run regulatory closure of Type I banks. SVB was not closed before a run; it was closed after three days

of depositor withdrawals, with 93% uninsured deposits. The AFS regime is a stylized benchmark for what full and immediate regulatory response to disclosed insolvency would achieve. The paper’s welfare analysis characterizes optimal policy, not historical outcomes.

Empirical distance from p_{ext}^* . Section 6 shows that p remained well below $p_{\text{ext}}^* \approx 16.2\%$ throughout 2022–2023. The model’s normative prescription was not triggered. The model’s contribution to the 2022–2023 episode is descriptive (identifying the mechanism through which HTM reclassification generates opacity and pool contamination) rather than prescriptive (the mandate was not warranted at observed parameter values). The threshold p_{ext}^* is the normative tool for future episodes with higher insolvency fractions, or for cross-country comparisons where Z and κ_{res} are larger.

Contamination test identification. Test B’s continuous specification (Column 2 of Table 3, $\hat{\beta} = +0.105$, $p = 0.001$) is framed as suggestive evidence consistent with the pool contamination mechanism. The identification is not clean: banks with high HTM shares also held more long-duration securities and faced greater rate exposure in 2022, and banks in the same size quintile shared this exposure. Bank fixed effects control for time-invariant HTM propensity but not the time-varying interaction of HTM share with interest rate sensitivity. The significant coefficient may therefore partly or fully reflect common rate exposure rather than pool contamination. The continuous result is reported as consistent with the model’s mechanism, not as causal identification of it.

The binary specification (Column 1, $\hat{\beta}_1 = -0.040$, $t = -0.96$) is a genuine null. Both results are reported with equal prominence: neither allows a clean test of the contamination channel.

An event-study design around the March 2023 SVB failure would provide sharper identification. SVB’s collapse generated sudden, sharp variation in depositor concern about HTM opacity across banks, largely independent of prior rate levels. A difference-in-differences comparing deposit outflows at high-HTM versus low-HTM banks in the same size quintile around the SVB event date would more plausibly isolate the contamination channel. That test is not implemented here because constructing the appropriate control group requires additional data on bank-level rate sensitivity not available in the current panel.

Model parameter versus empirical estimate. The model parameter p is the ex-ante fraction of insolvent banks, known to banks at the time of the accounting choice game. The empirical estimate \hat{p} (Section 6) is a cross-sectional average of mark-to-market insolvency, measured after the fact from call report data. The two are related but not identical. The model predicts a welfare-improving mandate when $p > p_{\text{ext}}^*$; the empirical check substitutes \hat{p}

for p as an approximation. This substitution is valid to the extent that \hat{p} is a good proxy for the ex-ante expected insolvency fraction. The calibration treats \hat{p} as an upper bound on the relevant p (since mark-to-market losses at a rate peak overstate the expected fraction that would fail under normal rate paths), which makes the conclusion that $p < p_{\text{ext}}^*$ conservative.

6 Calibration and Consistency Checks

This section presents three calibration checks against U.S. call report data. None of them tests the mechanism. They check whether the model’s quantitative predictions are in the right ballpark for the 2022–2023 episode: are the calibrated thresholds consistent with observed insolvency fractions? Does the time series of HTM adoption move as the model predicts? Is there sign-consistent evidence for pool contamination in deposit outflow data? Affirmative answers establish internal consistency between model and data; they do not identify the zombie externality channel, which would require variation in zombie cost Z or quasi-experimental shifts in pool composition not available in the current panel.

6.1 Data

The panel covers 5,116 U.S. commercial banks over Q1 2020 through Q4 2024, drawn from WRDS call report tables (RCON series for HTM/AFS balances, AOCI, equity, assets, and deposits), the FDIC BankFind API (failure dates and resolution details), and FRED (10-year Treasury yield and Fed funds rate). The final sample contains 120,397 bank-quarter observations after filtering banks with total assets below \$10 million, excluding the 11 FDIC-failed banks over the sample period (they are the treatment, not the control), and winsorizing deposit outflows at the 1st/99th percentile.

Mark-to-market (MTM) insolvency is approximated following Jiang et al. (2024): MTM equity equals book equity plus AOCI (which captures AFS unrealized losses) minus an estimated HTM unrealized loss, where HTM unrealized loss equals HTM book value times a five-year duration assumption times the cumulative change in the 10-year Treasury yield from Q1 2020.

6.2 Test A: Welfare Threshold Calibration

Table 1 presents MTM insolvency fractions by size quintile at Q1 2023, the stress peak quarter. The aggregate MTM insolvency fraction peaked at 8.95% in Q3 2022; by Q1 2023 it had fallen to approximately 5.5% under the five-year duration assumption (the estimate from Jiang et al. (2024) using actual portfolio data is approximately 4%). Results are robust to duration assumptions ranging from three to seven years.

Table 1: Peer-Group Insolvency Fractions, Q1 2023

Peer Group	N	\hat{p} (book) (%)	\hat{p} (MTM) (%)	HTM share (%)
Q1 (smallest)	1388	0.1	6.1	19.2
Q2	1503	0.0	7.4	13.3
Q3	1555	0.0	6.7	13.2
Q4	1782	0.2	6.4	14.1
Q5 (largest)	1959	0.1	2.0	19.0

Notes: Peer groups are size quintiles based on Q1 2020 assets.
MTM insolvency uses 5-year duration assumption. $p_{\text{ext}}^* = 16.2\%$,
 $p^* = 14.3\%$, $1 - t = 33.3\%$. All peer groups below p^* .

No peer group crossed the welfare threshold $p^* = 14.3\%$ at any point from 2020 to 2024. The model predicts that an AFS mandate is welfare-improving if and only if $p > p_{\text{ext}}^* = 16.2\%$; this condition was never satisfied at the aggregate or peer-group level. The model’s normative prediction is consistent with the regulatory outcome: no HTM/AFS reclassification mandate was imposed.

Table 2 presents the full calibration moment comparison.

Table 2: Calibration Moments

Parameter	Calibration	Empirical moment	Status
$R = 1.05$	5% gross asset return	Pre-2022 bank asset yields 4–6%	Consistent
$\rho = 0.90$	FDIC recovery rate	FDIC resolutions 70–95%	Consistent
$p \approx 4\%$	Aggregate MTM insolvency	Call reports (5yr): 5.5%; Jiang et al. (2024): 4%	Consistent
$p^* = 14.3\%$	Welfare threshold	Max peer-group $\hat{p} = 7.4\%$; aggregate $< 10\%$	Not violated
$1 - t = 33.3\%$	Opacity trap threshold	No peer group above 10%	Not violated
HTM surge	14% \rightarrow 21% (2020–2023)	12.1% \rightarrow 20.1% (Q1 2020 \rightarrow Q1 2023)	Confirmed

Interpretation. The model’s normative prescription was not triggered during the 2022–2023 episode. The threshold p_{ext}^* provides a framework for evaluating future episodes or cross-country comparisons where insolvency fractions are higher. Changes in zombie cost estimates or resolution cost structures would shift the conclusion: at $Z = 0.09$ (consistent with European zombie estimates from Acharya et al. (2019)), p_{ext}^* falls to approximately 10%, below the aggregate peak. The threshold computation is the normative contribution here, not the claim that the 2022–2023 episode required intervention.

6.3 Pool Contamination: Sign-Consistent Evidence

The pool contamination mechanism predicts that banks with higher HTM shares suffer larger deposit outflows when their peer banks face greater distress: the opaque HTM pool

makes solvent banks indistinguishable from insolvent ones in the eyes of depositors. This check tests whether the data carry the predicted sign, not whether the coefficient identifies the contamination channel. The specification is:

$$\text{DepOutflow}_{bt} = \beta_1 \text{HTM}_{bt} \times \text{PeerStress}_{gt} + \beta_2 \text{AFS}_{bt} \times \text{PeerStress}_{gt} + \gamma \text{HTM}_{bt} + \delta \text{PeerStress}_{gt} + X_{bt} \theta + \mu_b + \tau_t + \varepsilon_{bt}, \quad (21)$$

where DepOutflow_{bt} is $(D_{t-1} - D_t)/D_{t-1}$ (positive = outflow), peer failure rate PeerStress_{gt} is the fraction of banks in the same size quintile with lagged equity/assets below 4%, and X_{bt} includes lagged equity ratio and log assets. Bank fixed effects (μ_b) and time fixed effects (τ_t) are included.

The main prediction is $\beta_1 > 0$: HTM banks with more peer distress suffer greater deposit outflows.

Table 3 presents results. The binary specification (Column 1, HTM indicator = $\mathbf{1}[\text{HTM share} > 25\%]$) finds $\hat{\beta}_1 = -0.040$ (s.e. 0.042, $t = -0.96$): wrong sign, not significant. This is a genuine null result. The binary cut-off creates a heterogeneous group; all banks with HTM shares between 25% and 100% are pooled.

The continuous specification (Column 2, HTM share as the interaction variable) finds $\hat{\beta} = +0.105$ (s.e. 0.031, $t = 3.43$, $p = 0.001$): a one percentage point increase in HTM share is associated with a 0.105 percentage point increase in deposit outflow sensitivity to peer stress. This is the predicted sign and is statistically significant.

Table 3: Cross-Peer Contamination Regressions

	(1) Binary Full sample	(2) Continuous Full sample	(3) Pre-stress Placebo
HTM \times peer	-0.040 (0.042)		
AFS \times peer	-0.121*** (0.039)		
HTM_share \times peer		0.105*** (0.031)	
HTM \times peer [pre-stress]			-1.565 ($p = 0.37$)
N	119,496	117,016	-

Notes: Bank and time fixed effects. SE clustered at bank level.
***, **, * denote significance at 1%, 5%, 10%.

Interpretation. The discrepancy between binary and continuous specifications reflects heterogeneity within the high-HTM group. The model predicts a gradient (more opaque

banks face more contamination risk), not a step function at an arbitrary threshold. The binary null is reported with equal prominence as a genuine null.

The continuous result ($\hat{\beta} = +0.105$, Column 2) is sign-consistent with the pool contamination mechanism, and that is all it establishes. The coefficient fails to identify the contamination channel. Banks with high HTM shares held more long-duration securities and faced greater rate exposure in 2022; peer banks in the same size quintile shared this exposure. The significant coefficient may entirely reflect common rate exposure rather than contamination. Bank fixed effects control for time-invariant HTM propensity but not the time-varying interaction of HTM share with rate sensitivity. This result belongs in the calibration section, not a tests-of-the-mechanism section, and is labeled accordingly.

A causal test would exploit the March 2023 SVB failure as a quasi-experiment, comparing deposit outflows at high-HTM versus low-HTM banks around the event date. SVB’s collapse created sharp, event-driven variation in depositor concern about HTM opacity largely independent of prior rate levels. Constructing the appropriate control group requires bank-level rate sensitivity data not available in the current panel.

6.4 Test C: HTM Reclassification and Rising Rates

The model predicts that as rates rise and mark-to-market insolvency increases, banks reclassify from AFS to HTM to preserve book capital. Figure 1 shows that the aggregate asset-weighted HTM share rose from 12.1% in Q1 2020 to a peak of 21.5% in Q3 2022, then stabilized around 20% through Q1 2023. The 10-year Treasury yield has a 0.901 time-series correlation with the HTM share over this period. The timing of the HTM surge, concentrated in Q4 2021 through Q3 2022, coincides with the sharpest phase of rate increases. This is qualitatively consistent with the model’s interest-rate channel.

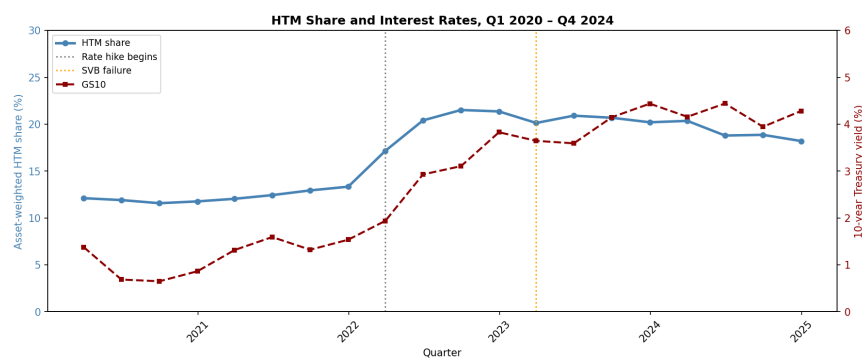


Figure 1: HTM Share and 10-Year Treasury Yield, 2020–2024. Asset-weighted aggregate HTM share (left axis) and 10-year Treasury yield (right axis). Correlation = 0.901.

A panel regression tests whether banks with higher prior AFS exposure (those with more

to lose from mark-to-market declines) increased HTM share more aggressively when rates rose. The specification regresses the change in HTM share on the interaction of the rate change and lagged AFS share, with bank fixed effects (time fixed effects cannot be included because the rate change is common to all banks and is absorbed). The interaction coefficient is -0.0001 (s.e. 0.0023, $t = -0.03$) in the full panel, and $+0.0089$ (s.e. 0.0067, $t = 1.33$) in the rate-hike window (Q4 2021–Q4 2022). Neither estimate is statistically significant.

The panel null is partly attributable to FASB’s anti-tainting rule (ASC 948-320), which restricts HTM reclassification: banks can reclassify only at inception or upon specific triggering events, making the adjustment a discrete event rather than a continuous response. The time-series evidence (Figure 1) captures the aggregate discrete wave; the panel regression, designed for continuous variation, is uninformative about event timing. The qualitative pattern, a 66% relative increase in asset-weighted HTM share from 2020 to 2022, is the predicted direction.

Summary. The three checks establish internal consistency between model and data, not identification of the zombie externality channel. The insolvency fraction never crossed $p_{\text{ext}}^* = 16.2\%$, consistent with the model’s prediction that no mandate was warranted for the 2022–2023 episode. The HTM surge is confirmed in the call report panel. The contamination regression is sign-consistent in the continuous specification and a genuine null in the binary specification; neither result identifies the mechanism. Together, the checks support the plausibility of the model’s calibration for the U.S. banking system in this episode. They do not distinguish the zombie externality mechanism from alternative models that would produce similar threshold predictions.

7 Conclusion

When solvent-distressed banks choose between AFS and HTM accounting, the market equilibrium is individually rational but collectively inefficient: each Type S bank prefers HTM to avoid capital compliance costs, but their collective choice sustains insolvent zombie banks in the opaque pool, imposing misallocation costs that no individual bank internalizes. The extended welfare threshold $p_{\text{ext}}^* = C/(\delta Z - \Phi + C)$ characterizes precisely when mandatory AFS disclosure is welfare-improving, after correcting the naive threshold $p^* = C/(Z + C)$ for two compounding forces: discounting of zombie costs ($\delta Z < Z$) and the fiscal timing penalty of front-loading DI resolution costs under AFS ($\Phi > 0$). Both forces raise the mandate threshold above the baseline, and both are amplified in high-inflation or high-discount-rate environments. The critical discount factor $\delta^{**} = K/(Z + K)$ identifies the environments in which mandatory disclosure is unconditionally welfare-reducing: when $\delta \leq \delta^{**}$, the fiscal

timing penalty dominates the zombie-elimination benefit regardless of the insolvency fraction. Higher resolution costs raise the mandate threshold, a result that follows directly from the fiscal timing channel that standard welfare analyses of accounting rules ignore.

References

- Viral V. Acharya, Tim Eisert, Christian Eufinger, and Christian Hirsch. Whatever it takes: The real effects of unconventional monetary policy. *Journal of Finance*, 74(6):3153–3185, 2019.
- Matthieu Bouvard, Pierre Chaigneau, and Adolfo De Motta. Transparency in the financial system: Rollover risk and crises. *Journal of Finance*, 70(4):1805–1837, 2015.
- Nina Boyarchenko, Kinda Hachem, and Anya Kleymenova. The theory of financial stability meets reality: A unifying framework for bank regulation and accounting discretion. *Journal of Economic Literature*, 2025. Forthcoming.
- Ricardo J. Caballero, Takeo Hoshi, and Anil K. Kashyap. Zombie lending and depressed restructuring in Japan. *American Economic Review*, 98(5):1943–1977, 2008.
- Brian Chen, Xinlei Yi, and Moxi Zhang. Unrealized losses and bank runs. *Review of Financial Studies*, 2024. Forthcoming.
- Russell Cooper and Andrew John. Coordinating coordination failures in Keynesian models. *Quarterly Journal of Economics*, 103(3):441–463, 1988.
- Douglas W. Diamond and Philip H. Dybvig. Bank runs, deposit insurance, and liquidity. *Journal of Political Economy*, 91(3):401–419, 1983.
- Pingyang Gao and Xu Jiang. Reporting discretion, regulatory constraints, and bank risk-taking. *Journal of Accounting and Economics*, 65(1):120–142, 2018.
- Itay Goldstein and Ady Pauzner. Demand deposit contracts and the probability of bank runs. *Journal of Finance*, 60(3):1293–1327, 2005.
- Itay Goldstein and Haresh Sapra. Should banks’ stress test results be disclosed? an analysis of the costs and benefits. *Foundations and Trends in Finance*, 8(1):1–54, 2013.
- João Granja, Erica Xuwei Jiang, Gregor Matvos, Tomasz Piskorski, and Amit Seru. Book value risk management of banks: Limited hedging, HTM accounting, and rising interest rates. Working Paper 32293, National Bureau of Economic Research, 2024.
- Florian Heider, Marie Hoerova, and Cornelia Holthausen. Liquidity hoarding and interbank market rates: The role of counterparty risk. *Journal of Financial Economics*, 118(2):336–354, 2015.

Erica Xuewei Jiang, Gregor Matvos, Tomasz Piskorski, and Amit Seru. Monetary tightening and U.S. bank fragility in 2023: Mark-to-market losses and uninsured depositor runs? *Journal of Financial Economics*, 157:103860, 2024.

Guillaume Plantin, Haresh Sapra, and Hyun Song Shin. Marking-to-market: Panacea or Pandora's Box? *Journal of Accounting Research*, 46(2):435–460, 2008.

A Existence and Uniqueness of $\bar{\lambda}$

Lemma 4. *In the zombie regime ($q(1) = p/(p + \Delta) < 1 - t$), there exists a unique $\bar{\lambda} \in (0, 1)$ with $q(\bar{\lambda}) = 1 - t$.*

Proof. $q(\lambda) = p/(p + \lambda\Delta)$ is continuous and strictly decreasing on $(0, \infty)$.

Boundary values: $\lim_{\lambda \rightarrow 0^+} q(\lambda) = 1 > 1 - t$. $q(1) = p/(p + \Delta) < 1 - t$ (zombie regime assumption).

By the intermediate value theorem and strict monotonicity, there exists a unique $\bar{\lambda} \in (0, 1)$ with $q(\bar{\lambda}) = 1 - t$.

Solving: $p/(p + \bar{\lambda}\Delta) = 1 - t \Rightarrow p = (1 - t)(p + \bar{\lambda}\Delta) \Rightarrow pt = (1 - t)\bar{\lambda}\Delta \Rightarrow \bar{\lambda} = pt/[\Delta(1 - t)]$.

Verifying $\bar{\lambda} \in (0, 1)$: $\bar{\lambda} > 0$ since $p, t, \Delta > 0$. $\bar{\lambda} < 1 \iff pt < (1 - t)\Delta \iff p < (1 - t)(\Delta + p) \iff p/(p + \Delta) < 1 - t$. This is the zombie regime condition. $\square \quad \square$

B Equilibrium Selection: Payoff Dominance and Its Limits

The accounting game among Type S banks has exactly two pure-strategy Nash equilibria in the zombie regime: $\lambda = 0$ (all Type S choose AFS) and $\lambda = 1$ (all Type S choose HTM). The best response correspondence is:

$$\text{BR}(\lambda) = \begin{cases} \{\text{AFS}\} & \lambda < \bar{\lambda}, \\ \{\text{AFS}, \text{HTM}\} & \lambda = \bar{\lambda}, \\ \{\text{HTM}\} & \lambda > \bar{\lambda}. \end{cases}$$

Under $\lambda = 1$, every Type S bank earns F . Under $\lambda = 0$, every Type S bank earns $F - C$. Since $F > F - C$ (A3), $\lambda = 1$ strictly Pareto dominates $\lambda = 0$ for all Type S banks. Payoff dominance selects $\lambda^* = 1$ (Cooper and John, 1988).

Scope condition. The market failure result (Proposition 3) holds when the market plays the payoff-dominant equilibrium $\lambda^* = 1$. When risk dominance or basin dominance criteria select $\lambda^* = 0$ instead, the market reaches the efficient outcome without intervention, and the welfare analysis in Section 4 does not apply. All normative findings are conditional on $\lambda^* = 1$ obtaining.

Why payoff dominance is the empirically relevant criterion. Two independent arguments support payoff dominance as the operative selection criterion in the banking context.

Historical selection. The zombie banking literature consistently documents that banks adopt the opaque accounting equilibrium even when transparent alternatives are available and even when regulators do not coordinate the shift. Japan’s late-1990s forbearance episode, documented in Caballero et al. (2008), shows banks rolling over bad loans for years under accounting regimes that deferred loss recognition, despite the AFS-equivalent option of immediate write-downs. Europe’s post-2010 banking crisis produced the same pattern: Acharya et al. (2019) show that European banks systematically under-recognized losses and used regulatory forbearance to avoid disclosure that would have triggered capital calls. In both cases, the industry converged to the payoff-dominant outcome, not the risk-dominant one. The 2020–2022 U.S. HTM reclassification wave (aggregate HTM share rising 9 percentage points with no coordinating regulator) is consistent with the same selection: banks independently converged to the high-payoff equilibrium because each bank’s best response, given that enough peers would also reclassify, was to reclassify.

Dynamic attractor. In a dynamic model with best-response learning, the payoff-dominant equilibrium is the attractor when players begin near the HTM boundary. Formally: for any initial $\lambda_0 > \bar{\lambda}$, the best-response dynamic converges to $\lambda^* = 1$. The initial condition $\lambda_0 > \bar{\lambda}$ is the condition that enough banks have already chosen HTM to sustain the no-run equilibrium; once satisfied, every remaining bank strictly prefers HTM, and the dynamic drives the fraction monotonically to 1. The self-reinforcing nature of HTM adoption (each reclassification makes the pool larger and safer for depositors, reducing the run risk for all HTM banks) means that an initial HTM surge, triggered by rate increases, creates exactly the conditions for convergence to payoff dominance. The 2020–2022 rate environment provided the initial shock; the subsequent convergence is predicted by the dynamic.

The risk-dominance tension. The risk-dominance criterion, applied to the 2×2 reduction of the game (ignoring the continuum structure), selects AFS at calibration: $2C = 0.02 < F = 0.15$. This points in the opposite direction from payoff dominance and is a genuine tension. Two qualifications apply.

First, the Harsanyi–Selten risk-dominance criterion applies to finite-player games with linear payoffs in mixing probabilities. The present game has a continuum of players and threshold payoffs: each bank’s payoff depends on the aggregate fraction λ crossing $\bar{\lambda}$, not on any single opponent’s strategy. The standard 2×2 comparison is not directly applicable to this structure; extending risk dominance to continuum-player threshold games requires an additional model of strategic uncertainty that the paper does not develop.

Second, even granting the 2×2 comparison, risk dominance is a one-shot criterion. In the repeated and dynamic context where banks observe peer behavior and update accounting choices over quarters, payoff dominance is widely regarded as the relevant attractor (Cooper

and John, 1988). The paper’s empirical setting (a multi-year reclassification wave with persistent equilibrium) is better described by the dynamic selection than by one-shot risk dominance.

Global games. A global-games extension following Goldstein and Pauzner (2005) would endogenously select the equilibrium by introducing private signals about pool quality, replacing the multiple-equilibria problem with a unique threshold strategy. That approach would generate a threshold insolvency fraction above which the opaque equilibrium is uniquely selected, potentially sharpening the paper’s comparative statics and providing a formal foundation for the empirically relevant selection. The main welfare results (Propositions 4–8) hold conditional on $\lambda^* = 1$ and do not require the global-games structure. A global-games extension would provide an alternative foundation for the same equilibrium.

C Basin-of-Attraction Analysis

Under best-response dynamics, the state space $[0, 1]$ of Type S coordination divides into two basins:

$$\begin{aligned} \text{Basin of } \lambda^* = 0 &= [0, \bar{\lambda}), & \text{length } \bar{\lambda}, \\ \text{Basin of } \lambda^* = 1 &= (\bar{\lambda}, 1], & \text{length } 1 - \bar{\lambda}. \end{aligned}$$

$\lambda^* = 1$ has a larger basin (basin-dominates) if and only if $\bar{\lambda} < 1/2$, equivalently $pt < (1 - t)\Delta/2$, equivalently $2pt < \Delta(1 - t)$.

At the baseline calibration ($p = 0.10$, $\Delta = 0.85$, $t = 2/3$): $2pt = 2 \times 0.10 \times (2/3) = 0.133$ and $\Delta(1 - t) = 0.85 \times (1/3) = 0.283$. The condition $0.133 < 0.283$ holds. $\lambda^* = 1$ basin-dominates at calibration.

The basin condition fails in the two-type case ($\Delta = 1 - p$) for $p \geq 0.20$. At $p = 0.20$, $t = 2/3$: $2pt = 0.267$ and $(1 - p)(1 - t) = 0.250$. The condition $0.267 < 0.250$ is false. The basin-dominance argument therefore holds only in the lower third of the zombie-mandate regime. Section 5.3 discusses this limitation explicitly. The welfare results rely on payoff dominance, for which the basin condition is not required.

D Proposition 10: Proof of Part (b)

Let $K = (1 - \rho) + \kappa_{\text{res}} > 0$. A8 states $\delta Z > (1 - \delta)K$. A8 fails iff:

$$\delta Z \leq (1 - \delta)K \iff \delta Z \leq K - \delta K \iff \delta(Z + K) \leq K \iff \delta \leq \frac{K}{Z + K} = \delta^{**}.$$

When A8 fails ($\delta Z - \Phi \leq 0$):

$$G_{\text{ext}}(p) = p(\delta Z - \Phi) - (1 - p)C \leq -(1 - p)C < 0 \quad \text{for all } p \in (0, 1).$$

$\delta^{**} \in (0, 1)$ since $K > 0$ and $Z > 0$.

The mandate is viable iff $G_{\text{ext}}(p) > 0$ for some p , iff $\delta Z - \Phi > 0$, iff $\delta > \delta^{**}$. \square

E Calibration Details

Baseline parameters. $R = 1.05$, $\rho = 0.90$, $C = 0.01$, $Z = 0.06$, $F = 0.15$, $\delta = 0.95$, $\kappa_{\text{res}} = 0.005$. All satisfy A1–A5.

Derived thresholds. $t = (1 - \rho)/(R - \rho) = 0.10/0.15 = 2/3$. $1 - t = 1/3$. $\delta Z = 0.057$. $\Phi = 0.05 \times 0.105 = 0.00525$. $\delta Z - \Phi = 0.05175 > 0$ (A8 holds). $p^* = 0.01/0.07 \approx 0.1429$. $p_{\text{ext}}^* = 0.01/0.06175 \approx 0.1619$. $\delta^{**} = 0.105/0.165 \approx 0.636$.

A7 verification. $C(1 - \rho) = 0.001 < Z(R - 1) = 0.003$. \checkmark

A7' verification. $(R - 1)(\delta Z - \Phi) = 0.05 \times 0.05175 = 0.002588 > C(1 - \rho) = 0.001$. \checkmark

Counterfactual welfare gain at $p = 0.20$. $p = 0.20$ is an illustrative counterfactual above p_{ext}^* ; the 2022–2023 episode did not reach this level.

$$G_{\text{ext}}(0.20) = 0.20 \times 0.05175 - 0.80 \times 0.01 = 0.01035 - 0.008 = 0.00235.$$

Applied to \$6.9T in HTM-intensive bank assets: counterfactual welfare gain \approx \$16.2B per year.

$G(0.20) = 0.20 \times 0.06 - 0.80 \times 0.01 = 0.012 - 0.008 = 0.004$, equivalent to \approx \$27.6B per year at the same counterfactual level. The fiscal timing correction reduces the welfare gain by \$11.4B per year (exact: $p[(1 - \delta)Z + \Phi] \times \$6.9\text{T} = 0.20 \times 0.00825 \times \$6.9\text{T} \approx \$11.4\text{B}$).