

Financing the Path to a New Currency Hedge

Capital, cash, and transition feasibility in Taiwan insurance

October 27 2026

Abstract

A financial institution can have adequate reported solvency and an attractive eventual portfolio yet lack the resources needed to implement a change. I define a conditional transition funding frontier: the minimum external equity required to reach a specified hedge target by a deadline, given a stress criterion, cash calendar and permitted repayable liquidity. A monthly linear program separates economic headroom from spendable cash and accounts for entry costs, post-stress exit, financing interest and principal repayment. Matched synthetic institutions with identical opening capital, cash and exposure require different funding solely because an existing asset becomes liquid at different dates. Allowing early asset monetization can eliminate that synthetic timing gap. A complete dated cash-flow ladder and financing menu contain the relevant information, so incremental behavioral explanatory power remains a proposed contribution. A joint asset–liability–forward reconciliation shows why avoided forward points cannot automatically be treated as economic savings. Eighteen cash, equity and FX-reserve reconciliations cover three reporting entities. An exact recurring-shock control calculation then re-solves policies with cash settlements, trading costs and adjustment delays. Earlier analytical, historical-market and stress experiments remain available and separately interpreted. Public statements do not identify institution-specific economic transition gaps, and no causal or incremental predictive validation is claimed. The result is an implementable theory and measurement framework with explicit empirical rejection conditions.

Keywords Currency mismatch; insurance; capital regulation; financial reform; foreign exchange hedging

JEL classifications G22; G28; F31; E58

1 Introduction

A favorable destination does not ensure a financeable path. Currency-hedge adjustment can change payments and contingent losses before an institution accumulates the earnings expected to finance it. Implementation depends on when resources arrive, whether they absorb losses or cover payments, and which legal entity can use them.

This paper defines a conditional transition funding frontier: the minimum external equity needed to attain a hedge target by a deadline at each available liquidity limit. Equity contributes both cash and loss absorption. Borrowing supplies cash and creates an obligation. An asset maturity or sale converts an existing claim into cash without automatically creating profit. These identities make the financing menu part of the measure rather than an afterthought.

A matched synthetic example initially gives different funding needs to institutions with equal opening capital, exposure, cash and near-term liquidity. Delaying the same receipt raises required equity from zero to 0.61% of the exposed block. The expanded-access test narrows that result: allowing early asset sales or suitable secured financing can erase the timing gap. A maturity association alone therefore cannot establish a financing obstruction. A complete cash-flow ladder and monetization menu contain the relevant information, so incremental predictive usefulness remains a proposed contribution.

Taiwan insurance provides a consequential measurement setting. Eighteen accounting bridges reconcile cash, equity and FX reserves across six reporting periods for KGI Life, Cathay Life and Fubon Life. A further derivative audit distinguishes nominal quantities, carrying amounts and maturity disclosures. Fubon’s reported aggregate nominal amount rises while its derivative-liability maturity sum falls. These objects cannot be substituted for an economic hedge ratio or a settlement calendar. The public accounts still do not identify stress-consistent liability values, marginal retained resource effects or the full legal financing set.

The economic benefit of adjustment needs separate verification. A joint asset–liability–forward reconciliation shows why avoided forward points are not automatically an economic saving. A newly solved recurring-risk model distinguishes noncash asset marks from cash hedge settlement and shows that capital scarcity and cash scarcity can push feasible hedge choices in opposite directions. Its earnings and resolution assumptions are stated and tested, not estimated as insurer primitives.

All earlier analytical, historical-market and stress experiments remain. They establish conditional results, contextual scale and failure conditions rather than actual institution-specific gaps. The fixed primary behavioral test concerns future common-equity cash issuance on a disclosed-target and issuance-access subsample. It cannot be executed honestly with the missing inputs. The paper is therefore a theory and measurement working paper; incremental prediction and credible causal identification remain requirements for a stronger publication claim.

2 The contribution: a conditional funding frontier

2.1 Mechanism comparison, page 1 of 2

Froot, Scharfstein, and Stein (1993; working paper 1992) make hedging valuable when external finance is costly relative to internal funds: preserving resources in states with valuable investment opportunities can prevent costly financing or forgone investment. The covariance between operating cash flows and investment opportunities matters. Their argument is about coordinating financial and real decisions, rather than minimizing measured earnings variance. It already gives a reason for the timing and state contingency of resources to matter.

Rampini, Sufi, and Viswanathan (2014) make financing and risk management compete for collateral. Their model is dynamic; scarce net worth can make investment or avoiding contraction more valuable than hedging. Consequently, constrained firms can hedge less. A claim that previous work ignores dynamics or that financial constraints necessarily increase hedging would misstate this benchmark.

The object developed here is different in purpose: the minimum external resources needed to execute a specified change in an existing portfolio by a specified deadline while remaining within separate capital and cash constraints. The portfolio destination is held fixed when measuring feasibility. Optimizing that destination is a separate question. “Transition gap” denotes this conditional implementation requirement, not a newly discovered reason for hedging.

Three distinctions organize the proposed contribution. First, a favorable terminal balance sheet does not establish that intermediate payments can be made. Second, cash borrowed against an otherwise sound position is different from capital that can absorb a loss. Third, a resource at a subsidiary or parent is not available to the insurance entity unless a legally and operationally feasible transfer occurs before the relevant payment. A discounted total, consolidated equity figure, or current cash balance can obscure these distinctions.

This is an application and measurement proposal within the financing-and-risk-management literature. A sufficiently rich version of an existing dynamic financing model can incorporate the same constraints. The paper therefore does not claim that its primitives are unavailable to those frameworks. Nor does a new name for an optimized shortfall establish an independent economic factor. The empirical burden is to show that an explicitly calculated requirement usefully summarizes the interaction of existing constraints and predicts implementation better than reasonable competing summaries.

The separating comparison is consequently between observational specifications. Consider two institutions matched on opening assets, currency exposure, reported capital, economic headroom, cash and a 30-day liquidity measure. Give them the same future asset receipt, but put it in different months beyond that initial liquidity window. A static specification predicts no difference. The transition calculation can predict different feasible paths and financing needs. An existing model supplied with the full contractual calendar can also predict the difference. That last qualification is essential to the novelty claim.

2.2 Mechanism comparison, page 2 of 2

The proposed discriminating prediction is a timing interaction: delaying a resource receipt should matter most when it crosses a binding transition payment or loss-absorption threshold. Moving a receipt between two slack dates should have little or no effect. Permitting a genuinely repayable bridge should attenuate a cash-timing difference when terminal repayment is feasible; it should not remove an economic-capital deficiency. Making a book reserve immediately recognizable, with no change in economic resources or legal financing access, should leave the economic frontier unchanged. These are conditional model predictions, rather than generic predictions that weak institutions hedge less.

The accompanying matched example is designed to be defeated. A complete dated cash-flow ladder distinguishes its two institutions without a transition-gap label. If the gap adds no predictive value once that ladder and capital interactions are modeled flexibly, its defensible role is a compact scenario calculation for planning and supervision. That may be useful, but it is a narrower contribution than a new explanation of institutional behavior.

Conversely, incremental prediction over current cash, solvency and exposure would establish only that those coarse measures omit useful path information. It would not establish that the underlying financing mechanism differs fundamentally from existing dynamic models. A stronger contribution requires a contractual timing or transfer-access variation that affects the calculated requirement and subsequent implementation in the predicted direction, while leaving the relevant destination economics and competing channels sufficiently unchanged.

There is also a resource-identity restriction that strengthens the test. The measure must use retained *net* cash and changes in economic net worth consistent with the assets, liabilities and hedges. Treating all forward points as an avoidable deadweight expense artificially manufactures a self-financing transition. A theory that works only under that accounting treatment fails before its behavioral predictions are tested.

The formal result below is deliberately modest. Under the linear scenario specification, the minimum equity required as a function of an available loan limit is non-increasing, convex and piecewise linear. Its slope and binding-date shadow prices identify which financing restriction is expensive. These properties follow from linear programming and feasible-set inclusion; their mathematical ingredients are standard. Their research value lies in transparent implementation, an auditable mapping from disclosures, and falsifiable institutional comparisons.

The present revision establishes the accounting distinction, an implementable conditional measure, a matched feasibility counterexample, six reported-account reconciliations and a newly solved recurring-risk experiment. It does **not** establish incremental prediction of actual insurer financing or hedge decisions. The existing market evidence remains descriptive context. Publication should present this as a theory and measurement working paper until the required institution-level inputs and behavioral panel are observed.

2.3 Existing liquidity measures are the harder benchmark

The comparison must extend beyond hedging theory. Bai, Krishnamurthy and Weymuller (2018) construct a liquidity mismatch index linking asset market liquidity with liability funding liquidity and show predictive content for crisis outcomes and government borrowing. EIOPA’s 2021 insurance-liquidity methodology and its 2024 stress exercise already evaluate insurers’ liquidity under adverse conditions. Separate cash and solvency analysis is established supervisory practice; this paper cannot claim that distinction as new.

The proposed frontier conditions on a specified *change* in portfolio and minimizes the resources needed to execute it jointly with capital and payment constraints. That is a narrower estimand than a reference-portfolio liquidity shortfall, but sufficiently rich stress testing with management actions can represent it. The matched example does not establish equality of a contract-sensitive liquidity mismatch index: its weights may already respond to the different maturities. An institutional validation must calculate a defensible liquidity-mismatch or comparable contractual-liquidity benchmark, rather than assume that two equal current cash ratios imply equal liquidity risk.

Accordingly, a useful outcome could be a transparent implementation of transition-constrained stress testing, without a new primitive mechanism or an empirically superior standalone statistic. Incremental prediction against these harder measures remains unestablished. The institutional test below explicitly includes such a benchmark alongside a flexible dated settlement ladder.

3 Economic accounting before transition financing

3.1 What is, and is not, an economic saving

Let S_t be domestic currency per unit of foreign currency, F_{0T} a forward rate, A^* the foreign principal, and h the hedged share of terminal foreign proceeds. In a one-period illustration with foreign return r_f and a fixed domestic liability L_T , the combined terminal resource balance is

$$W_T = A^*(1 + r_f)S_T + hA^*(1 + r_f)(F_{0T} - S_T) - L_T - \tau(h) - f(h). \quad (\text{G1})$$

Here τ is transaction expense and f the funding expense of collateral or settlement. Posted collateral principal remains an asset if recoverable; it is a cash commitment, not automatically an expense. Changes in the value of collateral must be accounted for separately. A bond’s credit losses, duration marks and any economic revaluation of insurance liabilities belong in the full calculation even though they are held fixed in this illustration.

With matching maturities and frictionless covered interest parity, $F_{0T}/S_0 = (1 + r_d)/(1 + r_f)$. The interest differential in forward points offsets the foreign cash return under a complete hedge. Removing that hedge changes exposure to S_T ; it does not simply remove a fee while preserving the same payoff. Cross-currency basis, dealer costs and collateral funding can create additional wedges, but each requires its own measurement. An observed forward quote cannot by itself separate them into causal or welfare components.

Table G1 evaluates assets, liabilities and the forward together. A domestic investment of 100 funds a foreign asset yielding 5%; the domestic liability promises 102. Set $S_0 = 1$, $r_d = 2\%$, and $F = 1.02/1.05$. The complete hedge covers the foreign terminal payment of 105, not merely its principal. Assumed

transaction and collateral-funding expenses total 0.25. These rates, expenses and probabilities are constructed examples, not estimates for Taiwan.

Domestic value at terminal date	Home appreciation	Unchanged spot	Home depreciation
Spot, domestic per foreign currency	0.90	1.00	1.10
Foreign asset payment	94.50	105.00	115.50
Domestic liability payment	-102.00	-102.00	-102.00
Complete forward payoff	7.50	-3.00	-13.50
Hedged net, after 0.25 expense	-0.25	-0.25	-0.25
Unhedged net	-7.50	3.00	13.50

Table G1. Joint asset–liability–hedge reconciliation. Expense principal is not counted twice as collateral and loss. Source: exact arithmetic in `src/hedge_accounting.py`.

Under physical probabilities (0.2, 0.6, 0.2), the unhedged expected net payment is 3 and the hedged payment is -0.25. The apparent expected improvement of 3.25 consists of 3 from changing the currency payoff and 0.25 from the assumed expenses. Positive pricing probabilities (17/35, 11/35, 7/35) value the foreign asset at 100 and the forward at zero after discounting at 2%. Under those prices, the unhedged net present value is zero and the hedged net present value is -0.2451. Thus the 3-unit expected carry advantage is not an additional 3 units of market value. Removing the specified expenses has value in this frictionless comparison; added distress, liability or financing costs could overturn that advantage.

The currency payoff is borne first by the institution’s residual claimants and loss-absorbing resources. Policyholders and creditors can become exposed when those resources are exhausted; a guarantee can transfer losses further. Higher expected accounting earnings, more immediately spendable cash, higher shareholder value and greater social welfare are distinct estimands. This paper’s conditional feasibility calculations do not rank them as equivalent.

The earlier continuous-time benchmark uses a retained flow $\rho c x$. Its formulas remain correct conditional on that resource flow. The application now requires a stricter interpretation: c must be a specified net retained-resource effect under the selected valuation and cash convention, or explicitly a hypothetical expense flow. The paper has not measured a positive net economic saving for any sampled insurer. In particular, none of the following experiments validates the claim by inserting a reported aggregate “hedging cost” directly into c .

4 Measuring the transition funding gap

4.1 A definition another researcher can implement

Specify an insurance legal entity, valuation date, eligible exposure denominator Q_0 , initial economic headroom K_0 , spendable unrestricted cash C_0 , a target $\bar{x}$, deadline T , and a set of admissible dated policies Π . The initial headroom excludes resources already needed to support unchanged liabilities and risks. A policy contains the hedge path, asset or liability changes if permitted, financing draws and legally available transfers. Its actions must use information available at each decision date. Specify a stress system $\mathcal{S}$, liability valuation, settlement calendar, adjustment delays, exit rule, credit covenants and permitted financing channels before calculating a gap.

Equity E adds both spendable cash and loss absorption. A loan facility with committed limit $\bar{L}$ permits draws but creates a matching obligation; it does not itself add economic capital. Let $\mathcal{F}(T, \bar{x}, \mathcal{S}, \Pi)$ contain the nonnegative pairs $(E, \bar{L})$ for which an admissible policy reaches the target by T , pays financing obligations and satisfies all specified capital and cash constraints. The conditional equity funding gap is

$$G_E(\bar{L}; T, \bar{x}, \mathcal{S}, \Pi) = \inf\{E \geq 0 : (E, \bar{L}) \in \mathcal{F}\}. \quad (\text{G2})$$

If no policy is feasible, the gap is infinite. Report the pair or frontier when financing instruments differ; a single number requires an agreed loan limit, financing prices and instrument set. Adding a stand-alone “capital gap” to a stand-alone “cash gap” double-counts equity that can satisfy both. Likewise, parent and subsidiary cash cannot be summed unless an admissible dated transfer supplies it to the relevant entity. A source of capital that is restricted from paying the modeled obligations needs its own state and cannot be treated as ordinary E .

The stress criterion is part of the object. A pointwise ability to absorb a prescribed shock at each date differs from survival of all recurrent shock sequences. A probability-of-breach constraint differs from either. Report scenarios as scenarios; there is no assumption-free transition gap hidden in a published balance sheet.

4.2 The implemented monthly linear program

Normalize $Q_0 = 1$ and let $x_t \in [0, 1]$ be the unhedged share chosen at the start of month t . Let $u_t \geq |x_t - x_{t-1}|$, $x_0 = 0$, a trading fee κu_t , retained *net* earnings $vx_t/12$, and a single time-zero loan draw $0 \leq L \leq \bar{L}$ repaid at month N . Earnings arrive at month end; they cannot finance that month’s initial trade. An existing asset of amount M becomes cash at the start of month m . The claim cannot be sold or pledged before its maturity under this selected financing set; allowing either channel defines a different frontier. It was already owned: receipt increases cash and reduces the noncash asset by the same amount, without adding economic capital. Before monthly income,

$$K_t^- = K_0 + E + \frac{v}{12} \sum_{j < t} x_j - \kappa \sum_{j \leq t} u_j, \quad C_t^- = C_0 + E + L + M \mathbf{1}\{t \geq m\} + \frac{v}{12} \sum_{j < t} x_j - \kappa \sum_{j \leq t} u_j. \quad (\text{G3})$$

End-of-month states add that month’s retained income. At N , loan interest $r_L LN/12$ reduces both capital and cash, and principal repayment reduces cash only. Require both pre-income and post-income capital to exceed $[a + \kappa(1 - a)]x_t$: a proportional shock a plus the cost of closing the remaining marked exposure, following the earlier exit-cost result. Require cash to exceed ℓx_t at both checkpoints, including after final loan repayment. The coefficient ℓ is a stipulated stress cash reserve. It is not an inferred regulatory liquidity rule or an observed insurer margin schedule. A different settlement design can make required cash increase with the *hedged* share instead; the recurring model below does precisely that.

Minimize E subject to these linear constraints, target $x_N = 1$, and any decision-calendar restrictions. An optional lexicographic second optimization holds minimum equity fixed and maximizes cumulative adjustment, with a small turnover/draw tie penalty. This selects a reproducible path from otherwise

multiple minimum-equity solutions. It is not a claim that the institution maximizes that secondary objective. The code exports binding constraints, dual shadow prices, dated paths, and direct accounting checks reconstructed independently from the constraint matrix.

Proposition G1 (funding frontier). In this linear specification, the conditional equity gap is non-increasing, convex and piecewise linear in the available loan limit on its feasible domain. A feasible loan can be left undrawn, so increasing the limit cannot worsen the optimum. Convex combinations of feasible policies and resource pairs remain feasible; the projection into resource space is a polyhedron, establishing convexity and piecewise linearity. These are standard linear-program properties. If an intervention changes only reported recognition and leaves the economic states, legal access and constraints unchanged, it leaves this frontier unchanged. If it changes distribution rights, collateral eligibility or funding access, that neutrality premise fails and the relevant constraint must change explicitly.

4.3 A decisive matched example and its limits

Both synthetic institutions begin with $K_0 = 8\%$, $C_0 = 0.5\%$, the same existing asset receipt $M = 3.5\%$, and identical assets, net currency exposure, book capital and opening solvency measures. Economic headroom and reported capital are distinct but each is held equal across the pair. The receipt is a zero-yield, fixed-value claim in this example; its discount rate is zero, and the 3% bridge rate is a financing spread. Changing its payment date therefore does not change opening asset value by construction. Introducing nonzero asset discounting would require matching present values explicitly. Their receipt falls outside the first 30 days in both cases, so opening cash and a liquidity measure restricted to that window also coincide. The target is complete unhedging within twelve months. Set $a = 10\%$, $\kappa = 0.25\%$, $\ell = 4\%$, annual retained net earnings $v = 4\%$, and annual loan interest 3%. These are deliberately favorable hypothetical earnings; they do not establish an attractive Taiwan reform destination.

Scenario	Minimum equity, % of exposure	Loan actually drawn, %	Deadline
Existing receipt in month 2	0.000	0.000	12 months
Same receipt in month 9	0.613	0.000	12 months
Month 9; loan limit 3%	0.000	2.746	12 months
Month 9; decisions every three months	1.078	0.000	12 months
Month 2; only 2% opening headroom, loan limit 20%	5.422	0.000	12 months
Month 2; zero retained net earnings	2.475	0.000	12 months

Table G2. Scenario-specific minimum equity. Cash and capital are checked before and after income; the loan is repaid at the deadline. The last two rows change their stated input and are not members of the otherwise-identical pair. Rounding is for presentation; machine-readable values and actual draws are supplied.

The timing wedge is not a lost terminal asset: both institutions own the same receipt. Early liquidity permits earlier adjustment and earlier earnings retention, which relaxes subsequent capital constraints. In the late case, restricted early cash limits the earnings-generating path, making initial equity useful even though it addresses both constraints together. A repayable bridge can restore that path. In the late case the minimum facility limit that permits zero additional equity is approximately 1.351% of

exposure. Table G2 reports a larger draw under a 3% limit because its secondary objective favors earlier adjustment; that draw is not the minimum bridge requirement. With only 2% headroom, even the maximum possible 4% annual retained flow cannot build the terminal loss-absorption requirement; a large loan limit does not repair that deficiency.




Figure 1. Conditional equity-liquidity frontier and maturity timing. Each point solves the monthly program. All inputs are synthetic. Curve shapes and shadow prices are properties of the stipulated instrument set, rather than estimates of Taiwanese financing supply.




Figure 2. Dated adjustment paths chosen by the documented secondary objective. Cash slack includes final repayment. These are selected feasible paths, not observed management decisions or uniquely optimal economic policies.

On an illustrative exposed block of NT\$1 trillion, 0.6125% corresponds to approximately NT\$6.13 billion of equity. This is a scale conversion, not an estimate of any insurer's requirement. The timing effect has no statistical confidence interval: a 108-case deterministic sensitivity grid varies net earnings, fee, cash reserve, timing and decision frequency. Its range is an assumption range. Among 54 paired early-versus-late configurations, 32 have a positive timing wedge and 22 have none; all 18 pairs with zero retained earnings have no wedge. The largest timing difference in this grid is about 0.900% of exposure. Earlier receipt cannot be worse under free cash retention in this specification, so the nonnegative sign is partly imposed by the admissible set. This is a diagnostic of when the interaction

matters, not empirical confirmation. If v is not measured, no narrow empirical uncertainty band is warranted.

The sharpest counterargument is accepted. With unrestricted access to an otherwise identical funding source, or when the affected dates are slack, the timing distinction can vanish. With a complete asset–liability settlement ladder, an ordinary cash-flow simulation sees the same difference. The measure’s added value must therefore be tested against that ladder, not merely against cash divided by total assets. A real-world approximation would match insurance legal entities on exposure and stressed economic headroom, then compare precommitted asset maturities or documented transfer access around a common adjustment deadline. The current disclosures do not establish such a matched natural experiment.

5 Institution-level accounting audit

5.1 Six reconciliations, with the perimeter explicit

The new dataset transcribes 84 core numeric fields, 30 reported operating items and 36 foreign-exchange reserve roll-forward inputs from six primary statements: 2024 and first-half 2025 for KGI Life, Cathay Life and Fubon Life. KGI is a stand-alone insurance company; the other two reports consolidate the life company and its subsidiaries. They are not financial-holding-company statements, but subsidiary resources in a consolidated life subgroup still need separate legal-access analysis. No parent cash balance is counted as freely transferable insurer cash.

For each period, two identities are reconstructed:

$$B_1 - B_0 = \text{profit} + \text{OCI} + \text{declared distributions} + \text{other disclosed equity movements} + \epsilon_B, \quad (\text{G4})$$

$$C_1 - C_0 = \text{operating cash} + \text{investing cash} + \text{financing cash} + \text{reported FX on cash} + \epsilon_C. \quad (\text{G5})$$

The six cash and six equity residuals are zero at the reports’ precision of NT\$1,000. Component-level page references, source URLs and source hashes are supplied. “Other equity” is reconstructed from named associate, subsidiary, noncontrolling-interest and share-compensation movements. It is not a residual plug. For example, a small Fubon restructuring item of NT\$44,000 is retained. Reclassifying earnings into stock capital or a special reserve is not fresh cash or new total equity.

Reporting entity, 2025 H1	Profit	OCI	Change in book equity	Change in cash
KGI Life, stand-alone	1.07	-29.08	-30.50	29.13
Cathay Life subgroup	19.01	-108.20	-89.67	247.40
Fubon Life subgroup	24.19	-108.73	-99.67	98.75

Table G3. NT\$ billion. Reported profit, OCI, equity and cash are different accounting objects. KGI: PDF pp. 10–12; Cathay: PDF pp. 12–15; Fubon: PDF pp. 9–11. Sources and 2024 opening/closing bridges accompany `data/institution_bridges.csv`.




Figure 3. First-half 2025 book-equity and cash bridges. Neutral bars are opening and closing stocks; positive and negative movements use common colors. Equity distributions are recognized declarations. Cash financing contains only cash actually reported as paid or received. “Other” comprises explicitly enumerated disclosed movements.

These bridges establish an accounting fact, not economic solvency: cash accumulation can coexist with an equity decline. The financing statements also distinguish bond proceeds from shareholder distributions. Fubon recognizes a cash dividend of approximately NT\$15.00 billion in the first-half equity statement without a corresponding dividend payment line in that period’s financing cash statement. KGI recognizes and reports paying NT\$2.50 billion. Neither a declared dividend nor a stock dividend is automatically cash already transferred to the parent at the same date.

The 2024 reported gross bond proceeds are NT\$10.00 billion for KGI, NT\$79.74 billion for Cathay’s life subgroup and NT\$59.98 billion for Fubon’s life subgroup. These are realized historical cash flows, not unused financing capacity or automatically admissible economic capital. Proceeds may fund other obligations, may have different capital eligibility, and do not establish willingness or permission to issue again. The illustrative NT\$6.13 billion conversion above is therefore comparable in units, but cannot be asserted to be affordable for any of these entities.

5.2 Reserve movements are a separate accounting bridge

Six additional reserve roll-forwards reconcile opening and closing FX valuation reserves through compulsory and additional provisions, recoveries and disclosed transfers. Their arithmetic residuals are zero. The net provisions also agree, with the opposite sign, with the income-statement reserve effect. The operating supplement records reported premiums, claims, interest income and FX gains or losses; it does not relabel them as directly observed cash settlement flows.

The first half of 2025 reveals a useful contrast. KGI records net reserve recovery of NT\$28.264 billion, with the reserve falling from NT\$30.705 billion to NT\$2.441 billion. Cathay records net provision of NT\$14.273 billion plus NT\$14.633 billion reclassified from other reserving accounts, reaching NT\$56.421 billion. Fubon records net provision of NT\$22.094 billion plus NT\$8.006 billion reclassified, reaching NT\$51.581 billion. Each movement is traced to its reserve note: KGI PDF p. 69, Cathay p. 89 and Fubon pp. 88–89.

These are different accounting responses in the same broad currency-stress period, not an identified explanation of hedge decisions. The disclosed transfers do not establish new spendable cash or an economic reduction in promises to policyholders. Likewise, an FX-gain income line is not the combined economic currency result: derivative effects, asset classification and overlay transfers can appear elsewhere. The reserve bridges make accounting substitution observable, while leaving the economic resource and settlement reconstruction open.

5.3 What remains unidentified

The data do not reconstruct instrument-by-instrument forward carry, cash hedge settlements, collateral timing, premium and claim payment calendars, stressed economic liability values, or legally available parent/subsidiary transfers. Operating cash flow is an aggregate that includes investment turnover and insurance activity; it is not cash generated by reducing hedges. Reported interest income and foreign-exchange gains are not substitutes for the missing marginal net-resource calculation. A reserve movement can affect reported profit without providing additional cash. KGI does not present a separate exchange-effect line in these cash statements; that convention is recorded, rather than interpreted as proof that foreign-currency cash has no exchange exposure.

The reported-number rounding bound is about NT\$5,000 per bridge under conservative summation of half-unit rounding errors. That bound concerns transcription and reporting precision only. It does not bound omitted economic liabilities, transfer restrictions or hedge terms. The current institutional transition gaps consequently remain *unidentified*, rather than estimated with narrow error bars. The missing-input table is part of the dataset. Full issuer PDFs remain available from the linked primary sources; the replication package supplies quantitative transcriptions and provenance rather than redistributing those documents.

The Insurance Bureau’s June 26, 2025 release reports a January–May life-insurance pre-tax loss of NT\$61.7 billion and end-May owners’ equity of NT\$2,034.5 billion. Those sector totals provide contextual aggregate checks, but not a mechanical sum check on three entities, consolidated subsidiaries or June-end accounts. The sample dates and perimeters differ. No unexplained sector residual is filled with a precise invented hedge estimate (Insurance Bureau 2025).

6 Re-solving adjustment under recurring losses and cash settlement

6.1 A small model with an economically distinct cash state

The historical replays later in the paper transport a first-event policy into a different environment. They remain useful stress diagnostics, but cannot identify the optimum in that environment. Here the policy is solved again. The horizon is twelve quarters, the action set is $x \in \{0, 0.1, \dots, 1\}$, and quarterly domestic-value increments on a unit foreign-currency block are $z \in \{-0.04, 0, 0.04\}$ with independent probabilities $(0.15, 0.70, 0.15)$. The increments are centered and additive, not a three-level mean-reverting spot process. There is no forecastable FX drift being treated as a free hedge saving. Interest differentials are set to zero in the FX settlement leg. A stipulated annual net resource effect of 2.5% is split into 2% retention and 0.5% owner distribution when fully unhedged. That positive net flow is an assumption, not inferred from the forward rate.

Asset currency marks are noncash. One-quarter hedges settle in cash. With $b = \kappa|x - x_{-1}|$ and quarterly retained resources $0.005x$,

$$K' = K - b + 0.005x + xz, \quad C' = C - b + 0.005x - (1 - x)z. \quad (\text{G6})$$

Trading and terminal closeout charges in this recurring experiment are quoted per initial eligible notional. The monthly pointwise exit reserve instead uses a post-stress marked-value convention; these are separate cost scenarios. The asset mark plus hedge settlement leaves capital exposed only to the unhedged share, whereas a foreign-currency appreciation can require cash payment on the hedge. This sign difference is the essential mechanism. It also supplies an independent identity, $K' - C' = K - C + z$, used to reduce the state space. Cash is already an asset component of net worth and is not added to capital a second time in the objective.

The objective discounts owner distributions at 3% annually and adds terminal residual capital after the cash-funded cost of returning the unhedged share to a hedge. A negative cash or capital state enters an absorbing resolution boundary with payoff $-0.10 - 5(-K)^+ - 5(-C)^+$. A benchmark attempting an unaffordable initial trade also pays a specified 0.10 penalty. These penalties are preference/continuation assumptions, not empirically measured bankruptcy costs or a complete shareholder limited-liability payoff. The objective is a transparent residual-value proxy; it is not a welfare optimum or an insurer default model.

Solve backward for each fee or decision delay. In the baseline, $\kappa = 0.0025$, $K_0 = 0.05$, $C_0 = 0.02$, and the opening position is fully hedged. The exact conditional action set is discrete; no claim of continuous-action optimality is made. The Bellman state records capital, previous action, quarter and cumulative currency increments. All reachable capital changes are multiples of 0.00025; thus no interpolation is required. A refinement to 21 nested actions on a 0.000125 capital lattice raises the exact baseline objective from 0.092980 to 0.093349 (about 0.40%) and leaves the initial share at 70%. This check does not establish a continuous-action optimum. The capital grid extends to 0.65. From the stated opening state even twelve consecutive positive increments plus maximum retained income cannot exceed 0.59 before fees, so its upper boundary cannot truncate a reachable baseline path.

An exploratory two-dimensional interpolation calculation proved materially sensitive near the resolution boundary. Its source and convergence diagnostics are retained, but its values are not used as the principal results. The exact lattice replaces that numerical approximation. Independent simulations

of 100,000 common paths reproduce each exact policy value within five Monte Carlo standard errors. Those intervals quantify simulation sampling only, not uncertainty about costs, cash-flow conventions, model probabilities or the value function.

6.2 Results and a sharp condition

The baseline re-solved initial unhedged share is 70%. The exact discounted objective is 0.09298 per unit exposure; the original first-event rule, discretized and evaluated in the same recurring environment, yields 0.05370. At the common opening state, that imported rule initially chooses 50%. The capital-only safe benchmark yields 0.05816. These comparisons concern this stated recurring model; they do not revise the original first-event theorem.

Policy and environment	Exact objective	Initial unhedged share	Model boundary crossings, %
Re-solved, 25 bp fee	0.09298	70%	0.725
Re-solved, no fee	0.10176	80%	0.428
Re-solved, 100 bp fee	0.06939	70%	6.541
Re-solved, decisions every two quarters	0.08664	70%	5.463
Original first-event rule transported	0.05370	50%	19.495
Capital-only safe rule	0.05816	50%	9.747
Always hedged	-0.09940	0%	60.592
Always unhedged	0.07431	100%	19.605

Table G4. Exact Bellman policy values; boundary frequencies from 100,000 simulated common paths. Boundary events are model cash/capital shortfalls, not forecasts of actual insolvency. Costs and adjustment restrictions change across the first four rows and policies are solved again. All other policy comparisons use the baseline environment.




Figure 4. Re-optimized and benchmark policies under recurrent shocks. Error bars describe Monte Carlo objective estimates; exact finite-state values are shown separately. These intervals must not be read as empirical uncertainty about Taiwanese insurers.

The fully hedged institution can cross the cash boundary because offsetting asset gains are not immediately spendable. Its high frequency is driven by the stipulated absence of additional liquidity facilities and asset sales. It is not evidence that real insurers should unhedge or that hedging is destabilizing in general. Permitting collateralized liquidity or asset monetization changes the admissible policy set and requires another solution.

The presumed resource flow is economically consequential even after re-optimization. With the annual net flow halved to 1.25%, the exact objective falls to 0.06149; with zero net flow it falls to 0.00926, and model boundary crossings rise to 25.04% on the common simulated paths. These cases are solved again, not evaluated with the baseline policy. At zero flow the initial unhedged share is still 70%, because cash settlement and capital risk remain in tension; that action is not evidence of a positive self-financing surplus. These results make validation of the retained net flow a prerequisite for policy interpretation.

One sharp implication can be obtained without numerical optimization. Ignoring trading fees, surviving either next-quarter 4% movement requires

$$\max\{0, (0.04 - C)/0.045\} \leq x \leq \min\{1, K/0.035\}. \quad (\text{G7})$$

Cash scarcity raises the lower feasible unhedged share; capital scarcity lowers its upper bound. At $K = 0.02, C = 0.005$, the bounds are approximately 0.778 and 0.571, so no action is robust to both movements. At $K = 0.05, C = 0.02$, the interval is approximately $[0.444, 1]$. A loan and an equity contribution relax different sides of this interval. Transaction fees tighten it, and inability to rebalance extends the relevant shock horizon. This conditional result explains why neither “less capital means less hedging” nor “more hedging is safer” is a universal prediction. Nine exact opening-resource cases are supplied in the replication results; their purpose is sensitivity, not estimation.

7 Testing asset access and the information in liquidity disclosures

7.1 A transition gap must survive an expanded financing menu

The matched example identifies a conditional timing obstruction. It does not establish that an insurer with marketable assets faces that obstruction. The initial calculation prohibited selling or pledging the existing receipt before maturity. This restriction is essential enough to test directly: expanding the financing menu can erase the result even though opening cash, capital and the eventual hedge destination stay fixed.

Allow a time-zero sale of an amount s of the existing domestic claim and a separate pledged amount p . The claim has already been included at fair value in opening economic capital. An execution discount h_s is therefore an *additional* loss relative to that value, not another recognition of a bond-price loss already included in K_0 . A sale generates $(1 - h_s)s$ of cash and reduces economic capital by $h_s s$. It also removes s from the later maturity receipt. Sale proceeds cannot be counted both now and at maturity.

For the secured option, the loan satisfies $L \leq (1 - h_p)p$, where h_p is a collateral haircut, and sale and pledge share the same asset constraint $s + p \leq M$. Pledging does not reduce capital by $h_p p$: the haircut limits borrowing capacity, while the asset is still owned. Interest remains a resource cost. If the pledged claim matures before the loan, its proceeds are locked until final repayment. The program releases those proceeds when it repays the debt. This prevents simultaneous use of the same asset as freely

spendable maturity cash and retained collateral. It is a specified secured contract, not an estimate of any insurer’s facilities.

Relative to equation (G3), the pre-income changes are

$$\Delta K_t^- = -h_s s, \quad \Delta C_t^- = (1 - h_s)s - (s + p)\mathbf{1}\{t \geq m\}, \quad s + p \leq M. \quad (\text{G8})$$

The baseline loan cash is already present in equation (G3). Final repayment releases matured pledged proceeds p , but a claim that has not matured cannot be converted into cash merely by releasing the lien. The code separately reconstructs these entries and checks every cash and capital checkpoint.

Proposition G2 (the monetization null). If the whole zero-yield claim can be sold immediately at its already recorded fair value, cash can be retained freely, and there are no other costs or restrictions, the minimum-equity gap is invariant to the claim’s maturity date within the horizon. Selling the whole claim gives the same initial cash $C_0 + M$ and leaves opening economic capital unchanged. Any policy requiring only later receipt can be financed from that initial cash by retaining it. Hence changing the receipt date cannot increase the minimum equity requirement. With optional costly sales or secured funding, enlarging the admissible set weakly reduces the minimum gap, but need not eliminate an underlying loss-absorption deficiency.

This null is more demanding than the original comparison with current cash. A researcher who observes an early-versus-late maturity association must also establish why the later claim cannot be sold or pledged on sufficiently favorable terms. Otherwise the association is not evidence of this particular transition mechanism.

	Minimum equity, % of exposure	Sale execution loss, % of exposure
Month-9 receipt; financing menu		
Original restricted menu	0.6125	0.0000
Full sale access; no execution loss	0.0000	0.0000
Full sale access; 20% execution discount	0.0131	0.5609
Full sale access; 40% execution discount	0.5035	0.8710
Committed secured line of 3%; collateral haircut 10%	0.0000	0.0000
Only 2% opening headroom; full sale access, no execution loss	5.4217	0.0000

Table G5. Synthetic minimum-equity calculations. All earlier monthly earnings, stress and trading-fee assumptions remain. The final row changes opening headroom and uses the month-2 receipt. A collateral haircut reduces credit capacity; it is not the sale execution loss in the final column. The secured loan incurs interest and must be repaid. The secondary path objective selects actual sale and pledge amounts from the minimum-equity set.

With access to the full 3.5% claim, the month-9 case requires no additional equity for execution discounts up to approximately 19.53%, under the deliberately favorable retained earnings in this example. This is not an estimate of bond bid–ask spreads or liquidation losses. It shows that the headline 0.6125% timing requirement is fragile to monetization in the example’s own parameterization. At a 20% execution discount, it falls to about 0.0131% of exposure, while realized execution loss is about 0.5609%. Even that smaller equity call cannot be called an economic benefit without valuing the loss and the resulting portfolio.

Minimum equity is consequently a feasibility statistic, not a preferred financing policy. Raising fairly priced equity transfers claims as it supplies resources; an execution loss destroys part of the existing asset value. Interest, issuance expenses, future income and duration changes also matter. Minimizing the equity call alone can select an expensive sale. A welfare or shareholder-value comparison requires another objective with explicit claimholder incidence. This distinction is especially important when a public bridge is compared with retained earnings or asset liquidation.




Figure 5. Expanded financing access and the synthetic minimum-equity requirement. The left panel solves the month-9 case across sale limits and execution discounts. The right panel shows the set of gaps consistent with a receipt dated somewhere in months 2–12, conditional on each financing menu. These are assumption ranges, not empirical confidence intervals. The calculation includes 486 sale cases, 24 secured-credit cases and 44 dated-calendar cases.

7.2 What coarse maturity buckets do and do not identify

Let Ω contain the exact calendars consistent with disclosed maturity bands, while holding the economic and financing inputs fixed. The institution may know its true contractual calendar even when the researcher sees only a band. Its conditional gap belongs to the interval with sharp lower and upper bounds

$$\left[\inf_{\omega \in \Omega} G_E(\omega), \sup_{\omega \in \Omega} G_E(\omega) \right]. \quad (\text{G9})$$

Not every interior value must be attainable on a discrete monthly grid. This is not a sampling interval. Nor is it generally the cost of a single policy that must work before the calendar becomes known. The latter requires nonanticipative policies and can exceed the worst fully informed scenario optimum. In this particular example, however, the uncertainty concerns only the arrival of a positive zero-yield claim. Earlier arrival can be ignored by holding cash. Therefore the latest feasible arrival is worst, the earliest is best, and the late-arrival policy can be used for all earlier calendars without anticipation. Endpoint calculations give sharp bounds in this restricted setting. The result does not extend automatically to uncertain claims, margin calls, currencies or state-dependent asset returns.

If the receipt can occur in months 2–12, the restricted-menu gap ranges from zero to approximately 0.7965% of exposure. Full costless sale access collapses the range to zero. In the stated example, full sale access at a 10% execution discount or the specified 3% secured line also yields zero additional equity throughout the calendar range. Thus finer timing data alone may be less useful than establishing whether an existing asset can be monetized. Conversely, a precise maturity date cannot compensate for an incorrectly measured economic headroom or marginal retained cash flow.

These are conditional bounds for a fully specified synthetic account. The actual insurers’ gaps are still not bounded by these numbers. Their complete stress-dependent cash obligations, net resource effects and admissible market access have not been reconstructed. Public maturity summaries should not be inserted as though they filled all those inputs.

7.3 A disclosure audit: three quantities that move differently

The new extraction records 49 derivative-liability maturity-bucket entries for December 2024 and June 2025 across the same three reporting entities. It preserves each report’s bands and instrument perimeter. KGI uses bands through 90 days, 91–180 days, 181 days to one year, and beyond one year. Cathay’s first band is less than six months; Fubon’s first band is less than one year. Uniform monthly payment schedules are not inferred from these different bands.

Reporting entity	December 2024 maturity-table sum	June 2025 maturity-table sum	June-2025 report PDF pages
KGI Life, FVTPL derivative liabilities	21.376	4.170	112–113
Cathay Life, listed swaps and forwards	90.395	12.489	144
Fubon Life, FVTPL and designated-hedge liabilities	42.293	16.536	137

Table G6. NT\$ billion. The sums describe the listed liability maturity tables, not the whole insurance payment calendar, gross derivative settlement legs or a comparable common-horizon liquidity requirement. Cathay’s listed instrument perimeter differs from its complete balance-sheet hedging category. Differences from selected book comparators are recorded rather than forced to zero.

For KGI, the two maturity-table sums equal the respective FVTPL liability carrying amounts exactly. That equality is a useful audit observation, not proof that every table uses fair values or that contractual outflows equal book values generally. The other reports explicitly describe contractual or estimated payment amounts, and their selected maturity totals differ from book comparators. No compliance judgment follows from these comparisons. The point is narrower: none of the extracted tables by itself supplies the paired gross exchanges, collateral triggers, currencies, settlement sequence and legal netting needed to calculate peak spendable-cash requirements.

As a conceptual accounting benchmark, AASB 7 Appendix B11D distinguishes net-settled swap payments from contractual exchanges where gross settlement applies; B11F highlights collateral and netting arrangements. The publicly available Australian compilation is used to check this distinction, not as a statement of the law governing Taiwanese insurers (AASB 2020). A zero net daily payment can coexist with a large temporary funding requirement if an outgoing leg must be paid before the incoming leg. An enforceable closeout-netting agreement on default also need not provide ordinary-course

payment netting. KGI's note expressly distinguishes settlement by agreement from the right to net on default (2025 half-year report, PDF p. 99).

Fubon's additional instrument table supplies a direct warning against treating liability values as hedge quantities. Its reported aggregate derivative nominal amount rises from NT\$1.4375 trillion to NT\$2.1337 trillion, approximately 48.4%, while the derivative-liability maturity sum falls by approximately 60.9%. Forward and swap composition also changes. Aggregate nominal amount is not a risk-adjusted hedge ratio: instruments can offset, reference different currencies or risks, and span different maturities. Nevertheless, the opposite movements reject using the liability maturity sum as a simple measure of aggregate nominal volume in this case.

Fubon derivative accounting object	December 2024	June 2025
Aggregate reported nominal amount	1,437.499	2,133.665
Gross derivative asset carrying amount	1.158	41.462
Gross derivative liability carrying amount	42.467	16.834
Net derivative carrying amount	-41.309	24.628
Liability maturity-table sum	42.293	16.536

Table G7. NT\$ billion. Nominal amounts and net marks: June-2025 report PDF p. 32. Gross assets are reconstructed from derivative and hedge-accounting details on pp. 25 and 28; gross liabilities use the corresponding reported liability categories. Gross assets minus gross liabilities reproduce each reported net amount with zero residual. Nominal principal is neither cash paid nor economic net exposure.

The 2025 net derivative asset mark is also not the cash already received during the half year. Unrealized valuation, realized settlement, collateral movements and contract replacement require different entries. Accounting designation likewise does not exhaust economic hedging: a derivative outside a designated hedge-accounting category can still be used to reduce risk. The issuer describes risk-management purposes separately from those categories. The extraction retains the classifications without treating them as a complete economic hedge ratio.

7.4 Resource access cannot be inferred from asset labels

KGI reports NT\$7.834 billion of government bonds and NT\$0.019 billion of bank cash pledged at June 2025. The deposits are recorded within refundable deposits. Its other-asset note separates an insurance-operation deposit of NT\$7.563 billion from derivative-trading deposits of NT\$0.271 billion and smaller items. The report also describes restrictions on returning the operation deposit while the business continues (PDF pp. 52–53 and 128). Those amounts are existing restricted positions, not an unused credit facility or a freely saleable buffer.

The derivative-trading deposit falls from NT\$14.244 billion at December 2024. That stock change should not be relabeled as a measured cash release: the pledged-asset note includes government bonds, cash and distinct deposit purposes, and valuation or composition changes can matter. Nor should pledged amounts recorded in other assets be subtracted a second time from the reported cash-and-cash-equivalents line. A calculation must map security ownership, pledged status, legal withdrawal rights, valuation and payment currency to the right ledger first.

These disclosures improve the measurement map without identifying a gap. They also supply an opposing institutional account: the issuers describe liquid securities, contingency arrangements and

adequate funding. A strong assertion that Taiwan’s insurers generally lack transition cash would have to confront that evidence. Accounting equity declines and positive derivative gains alone cannot establish a funding obstruction, while a management statement of sufficient liquidity is not a stress-calibrated capital estimate.

7.5 A sharper empirical rejection condition

Keep the primary outcome fixed at external common-equity cash issuance on the previously specified target-and-access subsample. Add monetizable unencumbered assets, execution-loss or borrowing-base assumptions, and legal settlement/transfer access to the ex ante input gate and to the benchmark. If inexpensive liquidation or committed secured credit makes the measured cash bottleneck disappear, the model cannot use that bottleneck to predict equity issuance. A remaining prediction must come from a separately demonstrated capital deficiency or another binding constraint.

A useful negative control follows from Proposition G2: among otherwise matched freely monetizable claims, maturity timing alone should not create the proposed financing effect. A timing coefficient that survives regardless of saleability, encumbrance or committed finance would weaken this mechanism and suggest omitted income, duration or selection effects. This prediction is more specific than a general association between illiquidity and financing, but its institutional test remains unexecuted. The completed contribution here is an auditable proxy-rejection example and an expanded feasible-set test, not evidence that the transition gap already outperforms conventional liquidity measures.

8 The empirical test that would establish the contribution

8.1 One primary prediction, chosen for measurability

The primary proposed outcome is next-twelve-month *external common-equity cash proceeds divided by lagged insurance-entity assets*, conditional on a documented hedge-adjustment target, a common deadline and continued access to equity issuance. Common equity is selected because its cash and loss-absorption effects match the instrument in the calculated equity gap. Bond issuance is not pooled with it unless the relevant admissible-capital specification is changed and verified. Stock dividends and reclassifications from reserves are excluded: they are not external cash. Gross equity issuance can be lumpy and rare, which limits statistical power and requires a longer entity-level sample than six reports. Debt proceeds, distributions, liability composition and actual hedge adjustment are secondary outcomes, not alternative searches for significance.

The signed prediction is positive on this stated access-and-target subsample: a larger ex ante equity requirement should predict greater external common-equity cash proceeds when equity is an available implementation margin. An institution without issuance access can instead miss the target, sell assets or shrink liabilities. Mixing those cases into an unconditional positive coefficient would test a different hypothesis. Decisions to announce a target and retain market access may themselves be endogenous; conditioning on them supports a restricted predictive claim, not causal identification.

The present six-report audit fails the input-identification gate: it supplies reported flows but not observed economic transition gaps. No gap regression is run and no significance claim is made. An absent cash-equity-issuance line in these statements must be distinguished from issuance outside the period, a parent issuance and a noncash share-capital increase. The saved protocol is retrospective and prospectively reusable; it is not preregistered.

For an identified panel, construct the gap using only filings, contracts and legal restrictions public before the decision window. Preserve publication timestamps and announcement dates. Compare held-out forecast error from a baseline using exposure, economic headroom, cash, profitability, duration mismatch, existing hedging, dated debt maturities a contract-sensitive liquidity-mismatch measure, and the settlement ladder with an otherwise identical model adding the conditional funding requirement. Exclude whole institutions and whole future periods from estimation. A complete dated ladder with flexible interactions is a harder and more relevant benchmark than current cash alone. Fit all transformations and tuning choices on training data only; report the full prespecified score comparison and uncertainty by institution and time rather than treating issuer-days as independent currency shocks.

Prediction remains different from causality. The BIS evidence that investor hedge flows helped move exchange rates in April–May 2025 rules out treating that episode mechanically as an exogenous currency instrument. A causal design would require a documented externally imposed transfer restriction, financing-access change or predetermined contractual timing variation, with an assessment of selection, anticipation and direct effects. No such exclusion restriction has been established in this paper. Timing coincidence, a significant market coefficient or the synthetic matched example does not supply one.

8.2 Tests intended to defeat the mechanism

Challenge	Prespecified weakening result
Simpler measures	Gap adds no held-out improvement beyond capital, cash and a flexible dated cash-flow ladder. Narrow the contribution to scenario compression.
Accounting relief	Results require adding a noncash recognition change to spendable cash or economic headroom. Reject that implementation.
Net benefit	Resource accumulation disappears after separating carry, fees, collateral funding and liability changes. Withdraw the self-financing application.
Episode and influence	Predictive improvement disappears outside May 2025 or after excluding one institution. Treat the finding as episode- or institution-specific.
Terms and transfers	The signed ranking reverses across plausible hedge contracts or legal-access bounds. Report an unidentified ranking.
Negative control	The gap predicts board turnover, used as an unrelated-action placebo, as strongly as the specified financing outcome after size and distress controls. Investigate omitted distress or governance exposure.
Timing mechanism	Moving an alleged binding receipt across the deadline does not change the calculated gap or behavior, while unrelated dates do. Reconsider the timing interpretation.
Instrument eligibility	Apparent equity proceeds are stock dividends, reserve transfers or a parent issuance outside the insurer. Reject that outcome measurement.

Table G8. Prospective rejection criteria. The board-turnover placebo could itself respond to distress, so it is a confounding diagnostic, not an automatic falsification. These institutional tests have not been executed without identified inputs. Completed model tests include the matched timing comparison, repayable bridge, capital-poor case, zero-income case, financing-frontier shape checks and 108-case sensitivity grid.

9 Policy incidence after accounting

The policy object is a constrained path, not a reward for eliminating a hedge. Equity adds genuine cash and loss absorption, with dilution or an expected return paid to investors. Retaining otherwise

distributable earnings can add both over time, at the expense of current shareholders’ distributions, but only when those earnings exist and are legally retainable. A committed loan adds payment capacity and a repayment obligation; interest reduces resources. An asset sale converts an existing claim into cash, potentially at a haircut, and can change duration or income. Liability changes can release liquidity while changing guarantees, surrender costs or policyholder protection. These interventions require different accounting entries and cannot be modeled as interchangeable increases in K_0 .

A book recognition change alone adds no economic resources in the neutrality benchmark. It can nevertheless matter if it changes binding distribution, collateral or financing rules; those effects should be modeled as changes in legal access, rather than asserted to be new cash. Public support can transfer risk from shareholders or creditors to taxpayers or a guarantee fund. The recurrent model’s penalties do not identify those transfers or their welfare weights, so it cannot rank public rescue policies.

The matched example supports a limited conditional policy conclusion: a properly priced, repayable bridge can remove a payment-timing obstruction when repayment is feasible. The low-capital and zero-income cases show why that conclusion does not extend to an underlying economic loss. Reporting the equity–liquidity frontier, its binding dates, identified input ranges and stakeholder incidence is more informative than labeling every institution “short of capital.” The empirical importance of those distinctions remains to be established with an identified insurer-level transition dataset.

10 Related literature and the scope of the contribution

Calvo and Reinhart (2002) document reluctance to allow exchange rates to move even under announced floating regimes. The present setting reverses a familiar balance-sheet exposure. An institution holding foreign currency assets against domestic currency promises can be harmed by appreciation of its home currency. This sign reversal does not require a new theory of exchange rate fear. It changes which institutions would benefit from accommodation and which direction of exchange rate movement creates their losses.

Du and Huber (2026) study foreign investors’ dollar security holdings and currency hedging across institutional investors. McGuire, Shim, Shin, and Sushko (2021) examine the connection between outward investment and dollar funding in emerging Asia. These studies establish the importance of gross portfolios and hedge arrangements. In particular, short-term hedges of long-term assets can generate funding and rollover problems. This paper takes the foreign asset block as given and asks how a change in its hedge can finance additional capital. It does not explain the initial portfolio allocation or estimate hedging demand.

Rampini, Sufi, and Viswanathan (2014) show how collateral constraints create a tradeoff between financing and risk management. Their finding that constrained firms may hedge less is central to interpreting any association between low capital and low hedge ratios. Such an association need not reflect confidence in a future rescue. In the continuous-time benchmark below, the financing link operates through retained expense savings rather than a collateral allocation between productive investment and hedging. This simplification permits a closed-form transition requirement, at the cost of omitting much of the underlying contracting problem.

Accounting can also alter actual behavior. Ellul, Jotikasthira, Lundblad, and Wang (2015) show that accounting treatments interacting with capital regulation affect insurers’ asset sales and gains trading. Sen (2023) documents that the regulatory recognition of risks affects life insurers’ hedging. Those

papers provide strong reasons to investigate a behavioral response to recognition rules. They do not justify treating every reduction in reported volatility as a change in economic capital. The neutrality benchmark below makes that distinction explicit.

Farhi and Tirole (2012) demonstrate how anticipated, imperfectly targeted support can make private risk choices strategic complements. The policy-response extension in this paper is a reduced-form illustration of that established concern. It is not a new collective moral hazard theorem and is not integrated into a general equilibrium with endogenous saving, export production, or monetary policy.

There is already directly relevant Taiwan research. Wang (2026) studies the 2026 foreign exchange accounting and reserve reforms with a dynamic asset-liability simulation framework. The repository abstract reports improvements in financial resilience under the combined mechanisms. This is a substantive opposing benchmark: reform can preserve resources and reduce measured shortfall risk. Only the repository record and abstract are used here; the underlying thesis has not been independently replicated. The present paper focuses on a different estimand, the resources required to reach a safer position while satisfying an explicit path constraint. A study of eventual or simulated average resilience need not answer that transition question.

Several results below are straightforward consequences of the model rather than deep mathematical innovations. The comparison argument is standard, and the aggregation result uses convexity of the positive part. Their value depends on whether the underlying transition constraint proves empirically relevant. The paper therefore treats its institutional mapping and falsification tests as part of the contribution, rather than relying on the presence of formal propositions alone.

11 Institutional motivation and measurement

11.1 The balance sheet and the reform sequence

The Central Bank of the Republic of China (Taiwan), or CBC, reports that life insurers held NT\$37.68 trillion in assets at the end of 2025, equivalent to 131.22 percent of GDP. Their foreign exchange hedging ratio was 50.23 percent and accumulated foreign exchange valuation reserves were NT\$613.7 billion. The average risk-based capital ratio was 315.09 percent, with two firms below the statutory 200 percent minimum. The top three insurers held 54.34 percent of sector assets (CBC 2026a, pp. 65–68). These figures establish scale and heterogeneity. They do not identify the model’s capital headroom.

Two interventions must be distinguished. First, the reporting rules provide for straight-line recognition of specified unrealized foreign exchange differences over the expected remaining life of eligible directly held amortized-cost debt instruments without designated foreign currency hedging. The provision applies from January 1, 2026, with recognition of the remaining amount upon derecognition. The official consolidated text inspected on October 26 also contains amendments scheduled for 2028; its article numbering should be matched to the historical version in any empirical coding (FSC 2026a, Article 12).

Second, the reserve rules amended on February 13, 2026 include reserve allocations and special surplus requirements linked to hedge reductions. They restrict cash distributions from specified surplus accounts and allow subsequent catch-up requirements or supervisory measures if earnings are insufficient. The rule’s designated hedge cost rate is 2.5 percent (FSC 2026b). That number is an administrative input. It is not an estimate of realized net savings, and it is not the retention parameter in the model.

The sequencing matters empirically. The end-2025 hedge ratio already precedes the 2026 accounting treatment and February reserve amendment. The U.S. Treasury’s July 2026 report cites a reported February 2026 hedge ratio of 45 percent (U.S. Treasury 2026, Box 1). Dates and definitions differ across these reports. Their juxtaposition is not an estimate of the reform’s effect and cannot establish that the rule caused the decline.

The sequence also includes material measures before 2026. An Executive Yuan Gazette order concerning foreign-exchange valuation reserves took effect on June 30, 2025. Issuer disclosures during the May 2025 stress describe reserve provisioning, use and reallocations. A treatment indicator that begins only in January 2026 would omit interventions already operating in the preceding year. A dated crosswalk of adoption and economic effects remains necessary; the earlier reserve order is included here as a confound, not treated as an isolated experiment (Executive Yuan Gazette 2025; Fubon Financial Holdings 2025).

11.2 Three quantities that must remain separate

Reported equity, regulatory eligible capital, and economic headroom answer different questions. Reported equity follows accounting recognition. Eligible capital follows the applicable solvency framework and transition provisions. Economic headroom is the amount available above a specified floor after all relevant asset and liability revaluations under a common stress. The paper’s state variable is the last quantity.

In particular, an aggregate risk-based capital ratio cannot be multiplied by foreign investments to obtain an economic exchange rate buffer. The denominator is risk based, the numerator may have regulatory adjustments, and the relevant exposure is net of liabilities and hedges. Reserve balances also require a reconciliation to the same valuation basis before they can be added to anything called capital. Double-counting an allocation already included in equity would falsely strengthen the model’s starting state.

For similar reasons, the ratio of foreign investments to total assets is not a net currency mismatch. Foreign currency policies, derivatives, foreign equity exposures, and other assets can change the sensitivity of the consolidated balance sheet. The legal hedge ratio has a defined denominator that excludes specified liabilities and assets. A researcher must reconstruct this denominator rather than infer net exposure from the foreign asset share.

11.3 The strongest institutional alternative

The FSC’s rationale emphasizes the long horizon of life insurance and the cost of repeatedly hedging long-lived assets. That rationale deserves to be taken seriously. An institution with stable domestic currency funding may be able to tolerate transient exchange rate variation more cheaply than one forced to refinance frequently. Accounting volatility can also trigger sales or funding consequences that destroy value. In that case, changing recognition can affect economic resources through behavior, even though it does not erase a currency loss.

The CBC has rejected the claim that it coordinated the accounting change to engineer depreciation. Its January 30 statement distinguishes offshore non-deliverable forwards from onshore swaps, describes different transmission channels, and rejects a predetermined appreciation barrier (CBC 2026b). The analysis here does not infer institutional intent or an exchange rate target. The narrower hypothesis is

that balance-sheet exposure can affect the cost of an exchange rate adjustment, whether or not anyone deliberately created that exposure.

12 A model of a self financed transition

12.1 Exposure and capital

Consider a block of net foreign currency exposure that is initially hedged. Its domestic currency value before a shock is normalized to one. This is an eligible block of a fixed portfolio, not the institution's entire balance sheet. Existing unhedged positions and other risks have already been allowed for in baseline capital.

At time t , the institution leaves a fraction $x_t \in [0, 1]$ of this block newly unhedged. Let $a \in (0, 1)$ be the stress decline in the domestic currency value of one unit of foreign exposure. The incremental stress loss is ax_t . A value of $a = 0.10$ means a 10 percent fall in the domestic currency price of foreign currency; it corresponds to an 11.11 percent increase in the reciprocal exchange rate. This convention avoids ambiguity about the denominator of an appreciation percentage.

Let k_t denote economic capital headroom above a required floor after the stress losses on the unchanged part of the balance sheet. All quantities are normalized by the same eligible block. Before the first currency shock, headroom follows

$$\dot{k}_t = g + \rho cx_t, \quad k_0 \geq 0. \quad (1)$$

Here $c > 0$ is the annual pre-shock cash saving per unit newly unhedged, $\rho \in (0, 1]$ is the proportion retained as additional economic resources, and g is the baseline retained economic surplus under the unchanged hedge. The baseline surplus is measured after fixed contractual costs, taxes, and the counterfactual distribution policy. It may be negative. Define $v = \rho c$ and $r = v/a$.

The retention parameter measures an increment relative to a specified cash-distribution counterfactual. A transfer from unrestricted retained earnings to a restricted account does not increase k at the transfer date. It can support a higher future ρ if the transfer prevents distributions that otherwise would occur. Similarly, actual cash retained before the reform cannot be counted again as reform-generated capital.

The sign assumption $c > 0$ concerns pre-shock cash flow. It does not assert that currency carry is a free expected return. Avoided forward expense can be compensation for exposure to exchange rate risk; the model represents that exposure through the adverse event. A structural asset-pricing model would need to price the entire state-contingent payoff. The present model instead asks a financing question conditional on a positive pre-event cash saving.

12.2 Timing and the safety criterion

The initial benchmark requires the institution to absorb the specified incremental currency stress at every pre-shock date:

$$k_t \geq ax_t, \quad k_t \geq 0. \quad (2)$$

Capital paths are absolutely continuous and controls are measurable. A policy is admissible on a horizon if it satisfies the accumulation equation and both inequalities throughout that horizon, up to changes on sets of measure zero. Exposure can be increased or reduced without transaction costs in the benchmark. A mandatory one-way reduction in the hedge would impose an additional control restriction, especially relevant to Proposition 6.

One interpretation is a first adverse currency event with Poisson intensity λ . The problem terminates at that event; the model does not recycle the buffer through repeated shocks. With discount rate δ , the pre-event savings objective over horizon H is

$$\max_x \int_0^H e^{-(\delta+\lambda)t} c x_t dt. \quad (3)$$

The main result holds for any strictly positive time weights. Constraint (2) is a prudential standard for one prescribed stress. It is not a complete definition of insurance solvency, a probability-zero-loss requirement under every possible shock, or a welfare theorem. Losses larger than a , changes in liabilities, and repeated events remain outside this benchmark.

12.3 Why use this small model

The model makes the financing sequence observable. Current headroom permits current exposure; exposure releases a cash flow; only retained cash raises subsequent headroom. It excludes endogenous asset sales, new business, liability repricing, equity issuance after time zero, and changes in the eligible exposure block. Those channels can relax the constraint, but they also require additional resources or contractual changes. The closed-form results should be read as requirements conditional on holding those channels fixed.

The economic opportunity cost of retaining capital is omitted from the benchmark. This makes the fastest safe transition an optimistic bound on what retained savings alone can accomplish. Adding a cost of retention can delay or reduce the desired transition. It cannot make a path that violates the accounting identity feasible without another source of funds.

13 Analytical results

13.1 The fastest safe path

Proposition 1. Among admissible policies starting from the same k_0 , the policy

$$x_t^* = \min\{1, k_t^*/a\}, \quad \dot{k}_t^* = g + v \min\{1, k_t^*/a\} \quad (4)$$

weakly dominates every other policy in both headroom and exposure at every date on which an admissible path exists. It therefore maximizes (3). For $k_0 < a$, before the first hit of zero or a ,

$$k_t^* = -\frac{g}{r} + \left(k_0 + \frac{g}{r}\right) e^{rt}. \quad (5)$$

The economic reason is specific to the assumptions. Increasing exposure up to the safe ceiling produces more current savings and never reduces future headroom before the shock. There is no investment opportunity or adjustment cost competing for the same retained cash. This monotonicity, rather than risk neutrality in a general insurance problem, produces the result. Appendix A gives the comparison proof.

13.2 A profitable destination can be unreachable

Proposition 2. Suppose $-v < g < 0$ and $0 \leq k_0 < a$. Define

$$k^{\text{crit}} = -\frac{g}{r} = \frac{-ga}{v}. \quad (6)$$

If $k_0 > k^{\text{crit}}$, full unhedging under the prescribed stress is reached in finite time

$$\tau_S = \frac{1}{r} \log \left(\frac{a - k^{\text{crit}}}{k_0 - k^{\text{crit}}} \right). \quad (7)$$

If $k_0 = k^{\text{crit}}$, the fastest safe path remains at that level and never reaches a . If $0 < k_0 < k^{\text{crit}}$, the maximal safe capital path reaches zero at

$$\tau_F = \frac{1}{r} \log \left(\frac{k^{\text{crit}}}{k^{\text{crit}} - k_0} \right), \quad (8)$$

and no policy is admissible indefinitely without another source of resources.

At full unhedging, the pre-shock surplus $g+v$ is positive by assumption. Nevertheless, a weak institution cannot assume enough exposure to generate that surplus safely. Its safe savings are too small to offset the baseline deficit, so its capital falls and its feasible exposure shrinks. Requiring retention is ineffective below this threshold unless another part of the arrangement changes.

For $g = 0$ and $0 < k_0 < a$, the transition takes

$$\tau_S = \frac{a}{v} \log \left(\frac{a}{k_0} \right). \quad (9)$$

A firm starting at zero cannot initiate the safe transition when $g = 0$. If $g > 0$, baseline surplus can supply the initial resources even from zero. If $g + v \leq 0$, a firm below a cannot reach a sustainably self-financing fully unhedged position through this mechanism. These cases show why omitting the baseline business can materially overstate a retention policy's effectiveness.

13.3 Capital required for a deadline

Proposition 3. Suppose $-v < g \leq 0$. To reach $k = a$ safely by a specified date $H > 0$, the minimum nonnegative up-front injection is

$$J_H(k_0) = [z_H - k_0]_+, \quad z_H = k^{\text{crit}} + (a - k^{\text{crit}})e^{-rH}, \quad (10)$$

where $[y]_+ = \max\{y, 0\}$ and $k^{\text{crit}} = 0$ when $g = 0$.

The deadline converts a qualitative capital concern into a resource requirement. The expression is a minimum over safe paths under the stated channels, rather than a forecast of an equity issuance. An initial injection exactly equal to $k^{\text{crit}} - k_0$ only reaches the stationary threshold when $g < 0$; a finite completion date requires strictly more. Equation (10) includes that additional amount.

The target z_H falls with the deadline and approaches k^{crit} as H tends to infinity. With no operating deficit, an arbitrarily long deadline makes the infimum injection zero. With a deficit, waiting cannot eliminate the threshold. The economic distinction is between a transition that is slow and one that cannot finance itself.

13.4 The distribution of capital matters

Consider institutions indexed by i , with eligible exposure weights $w_i > 0$ summing to one. Suppose transition parameters are common and that existing capital cannot be transferred between institutions. Up-front external injections may be assigned to each institution, but they must be nonnegative. The minimum sector budget is

$$\mathcal{J}_H = \sum_i w_i [z_H - k_{i0}]_+. \quad (11)$$

Proposition 4. At fixed mean headroom $\bar{k}_0 = \sum_i w_i k_{i0}$,

$$\mathcal{J}_H \geq [z_H - \bar{k}_0]_+. \quad (12)$$

A mean-preserving spread in initial headroom weakly increases $\mathcal{J}_H$. If capital can instead be transferred freely before the transition at no cost, the right side of (12) is attainable.

This result concerns the mobility of loss absorption. Strong institutions may hold more than the transition target, while weak institutions cannot borrow that prudential headroom without an actual transfer, guarantee, merger, or other enforceable arrangement. An industry average implicitly assumes away this restriction. Corporate groups may relax it, but fungibility must be demonstrated rather than presumed.

With different a_i, c_i, ρ_i, g_i , the institution-specific target z_{iH} follows from (10), and the budget is $\sum_i w_i [z_{iH} - k_{i0}]_+$. A mean-preserving spread in capital alone has the simple ordering in Proposition 4 only when the relevant targets are common, or when the joint distribution is otherwise restricted. This qualification matters if stronger firms also have cheaper hedges or more profitable liabilities.

13.5 What accounting recognition can and cannot do

First consider a recognition parameter that changes reported results but leaves actual cash flows, distributions, and economic constraint (2) unchanged. Equations (1)–(3) contain no such parameter. The economic feasible set and fastest safe path are consequently unchanged. This is an accounting-neutrality benchmark, not a claim that actual regulation has no behavioral effects.

Now consider a mechanical exposure rule that follows a recognition-sensitive ceiling:

$$x_\alpha(k) = \min\{1, k/(\alpha a)\}, \quad 0 < \alpha \leq 1; \quad x_0(k) = 1. \quad (13)$$

The parameter α is the fraction of the stress recognized in that behavioral ceiling. It is not the fraction of Taiwan's amortization schedule entering statutory capital. The rule is an illustrative policy, not a derived private optimum. True incremental stress shortfall remains

$$S_\alpha(k) = [ax_\alpha(k) - k]_+. \quad (14)$$

Proposition 5. For $0 < \alpha < 1$,

$$S_\alpha(k) = \begin{cases} k(1/\alpha - 1), & 0 \leq k < \alpha a, \\ a - k, & \alpha a \leq k < a, \\ 0, & k \geq a. \end{cases} \quad (15)$$

The largest shortfall over $k \geq 0$ is $a(1 - \alpha)$, attained at $k = \alpha a$.

For $g = 0$ and $0 < k_0 < \alpha a$, the time to full *reported* unhedging is

$$\tau_R = \frac{\alpha a}{v} \log \left(\frac{\alpha a}{k_0} \right). \quad (16)$$

At that date true headroom is only αa . If full unhedging is maintained, another $(1 - \alpha)a/v$ years is required to reach $k = a$.

The recognition rule can therefore accelerate the true accumulation of resources while opening a loss interval. It may reach $k = a$ earlier than the safe policy. Calling the entire effect illusory would be incorrect because cash savings are real by assumption. Calling the institution safe because reported unhedging is complete would also be incorrect because the loss-absorbing stock is still incomplete.

The peak in (15) is a comparative-static envelope, not necessarily a realized path maximum for every starting point. A firm beginning with $k_0 > \alpha a$ starts on the declining branch. The computational comparisons below account for starting capital explicitly.

14 Numerical experiments

14.1 Design and reproducibility

The experiments use a common illustrative case: $a = 0.10$, $c = 0.025$ per year, $\rho = 0.80$, and $k_0 = 0.02$, giving $v = 0.02$ and $r = 0.20$ per year. The discount rate is 0.03 and, where used, the constant first-event intensity is 0.10 per year. These values are chosen to make the financing sequence legible. In particular, the choice of $c = 0.025$ is not inferred from the FSC's designated rate and does not calibrate actual hedge cash flows. A year is the time unit; capital and losses are fractions of the eligible exposure block.

The replication package contains the complete source, machine-readable outputs, figure files, and an executable Python script. The script uses closed forms, numerical integration, a Bellman solver, and

Monte Carlo as separate methods where possible. Every published numerical table is generated from those outputs. There are 158 automated numerical checks in the reported run. They check internal consistency rather than external validity.

14.2 Experiment 1 transition feasibility

With $g = 0$, the fastest safe transition takes 8.047 years. Immediate full unhedging would accumulate enough true headroom in 4 years, but would start with an 8 percent stress shortfall. These are two different financing paths to the same headroom target.

Introducing a baseline deficit of $g = -0.005$ per year produces $k^{\text{crit}} = 0.025$. The original 2 percent starting headroom is then below the threshold. Even though the fully unhedged position has a positive 1.5 percent annual pre-shock surplus, a safe transition from that starting point is infeasible without additional resources. The maximal safe path exhausts headroom in 8.047 years. Starting exactly at 2.5 percent instead leaves the institution stationary. Figure 6 shows the transition paths and the divergence of completion times near the threshold.




Figure 6. All parameters are illustrative. The left panel uses a baseline deficit of 0.5 percent of the eligible exposure per year. Paths terminate at zero headroom. The dashed horizontal line is the full stress target and the dotted line the viability threshold. The right panel omits completion times above 30 years and infeasible starting points.

The code checks 24 parameter pairs against an independently integrated capital differential equation. Maximum errors are required to be below 10^{-8} in normalized capital. This confirms the implementation of the transition equations. It does not establish that a constant deficit or fixed stress describes Taiwan.

14.3 Experiment 2 recognition and the arrival of the shock

The recognition experiment follows rule (13) for four values of α . For a first shock with constant intensity, expected discounted shortfall is

$$L_\alpha = \int_0^\infty \lambda e^{-(\delta+\lambda)t} S_\alpha(k_t) dt. \quad (17)$$

The integral prices the time during which resources are insufficient for the specified stress. It is not the probability of an insurer default: k is headroom above a prudential floor, and the event model omits other balance-sheet items.

Recognized fraction α	Full reported unhedging in years	Full true stress capacity in years	Expected discounted shortfall as percent of eligible exposure
1.00	8.047	8.047	0.000
0.75	4.957	6.207	0.550
0.50	2.291	4.791	1.052
0.00	0.000	4.000	1.355

Table 1. Numerical quadrature of the first-event loss under the common illustrative parameters. Full true stress capacity means $k \geq a$ and not the absence of every possible insurance risk.




Figure 7. Left panel gives newly unhedged exposure; right panel gives true stress shortfall before the first shock. Lower recognition releases the full flow of savings earlier but permits exposure to outrun accumulated capital.

The result illustrates both the appeal and cost of lower recognition. The $\alpha = 0.5$ path reaches true stress capacity more than three years before the strict safe path. It also produces positive losses if the shock arrives during the transition. Neither completion time nor reported earnings alone orders these policies.

For immediate full unhedging, let $D = a - k_0$, $q = \delta + \lambda$, and $T_A = D/v$. Direct integration gives

$$L_A = \lambda \left[\frac{D}{q} - \frac{v}{q^2} (1 - e^{-qT_A}) \right]. \quad (18)$$

The code verifies this formula against quadrature and checks all four policies with 500,000 common first-event draws, using seed 20261026. The Monte Carlo standard error for the aggressive policy’s mean loss is 0.00331 percentage points of exposure. Simulation uncertainty is reported separately from assumptions about the event intensity.

14.4 Experiment 3 aggregate capital and the weakest institutions

With $g = 0$, a five-year deadline requires initial headroom of $z_5 = 0.036788$. Consider two equally sized institutions and mean headroom of 4 percent. If each has 4 percent, neither needs an injection to meet the deadline. If their headrooms are 1 and 7 percent, the weaker firm requires 2.6788 percent of its own eligible block. The sector needs 1.3394 percent of aggregate eligible exposure despite having exactly the same average capital.

This gap is generated by a restriction on transfers, not by a change in aggregate resources. Costless pooling would eliminate it in this example. A guarantee supplied by a stronger institution could also change the relevant constraint, but only if it provides enforceable loss absorption under the stress. Treating a group name or common ownership as proof of fungibility would bypass the economic issue.




Figure 8. The left panel shows 200 deterministic lognormal-quantile institutions at a common mean of 4 percent; the capital axis is logarithmic. The dashed threshold is the five-year target. The right panel extends the original dispersion experiment to transfers of surplus with alternative value-retention efficiencies. No actual institution sample or estimated transfer efficiency is represented.

The dispersion exercise does not rely on a favorable random draw. It uses fixed normal quantiles and a monotone increase in their log-dispersion, holding the mean exactly constant. The full grid is saved.

Each case is checked against the pooled-capital lower bound. The theorem supplies the ordering; the exercise quantifies its size for a transparent hypothetical distribution.

14.5 Experiment 4 optimizing when early losses are allowed

Strict stress safety is a policy choice. To evaluate an opposing perspective, set $g = 0$ and allow shortfalls, penalized at rate $\theta \geq 0$ per unit. The criterion is

$$V(k_0) = \sup_x \int_0^\infty e^{-qt} [cx_t - \theta\lambda(ax_t - k_t)_+] dt, \quad \dot{k}_t = vx_t. \quad (19)$$

This is a loss-tolerant policy objective, not measured aggregate welfare. The parameter θ can include disruption or policyholder costs beyond a unit of uncovered stress loss. It does not represent an estimated insurance guarantee subsidy. The problem continues to terminate at the first event and assigns zero continuation value there, including when no shortfall occurs. That boundary makes the comparison well defined but limits its interpretation.

For two fixed rules, the tradeoff can be expressed analytically. The aggressive rule obtains pre-event benefits $B_A = c/q$. The strict safe rule obtains

$$B_S = c \left[\frac{k_0}{a} \frac{e^{(r-q)\tau_S} - 1}{r - q} + \frac{e^{-q\tau_S}}{q} \right], \quad (20)$$

using τ_S from (9), with the first fraction replaced by τ_S when $r = q$. The aggressive rule beats the safe rule if and only if

$$\theta < \theta^* = \frac{B_A - B_S}{L_A}. \quad (21)$$

The illustrative cutoff is 5.218. Equation (21) compares those two rules only. An intermediate strategy may beat both.

The entire continuous-control problem can also be solved analytically. This produces a sharper result than a comparison of selected exposure rules.

Proposition 6. Suppose $g = 0$, $0 \leq k_0 < a$, $\lambda > 0$, and $r, q = \delta + \lambda > 0$. Under objective (19), if $\theta \leq c/(\lambda a)$, full unhedging is optimal at every capital level below a ; set $K_\theta = a$. For larger penalties, when $r \neq q$ and the bracket is positive, define

$$K_\theta = a \left[\frac{q + (r - q)\theta\lambda a/c}{r} \right]^{r/(q-r)}. \quad (21a)$$

When $r = q$, replace this expression with $K_\theta = a \exp(1 - \theta\lambda a/c)$. When $r < q$ and $\theta \geq cq/\{(q - r)\lambda a\}$, set $K_\theta = 0$. An optimal feedback policy in all these cases is

$$x^{\text{opt}}(k) = \begin{cases} 1, & 0 \leq k < K_\theta, \\ k/a, & K_\theta \leq k < a, \\ 1, & k \geq a. \end{cases} \quad (21b)$$

The value for $k \geq K_\theta$ is the safe-path value $B_S(k)$ from (20), extended continuously to zero. For $k < K_\theta$, put $T = (K_\theta - k)/v$; the value is

$$\int_0^T e^{-qt} [c - \theta\lambda(a - k - vt)] dt + e^{-qT} B_S(K_\theta). \quad (21c)$$

The proof in Appendix A verifies the Hamilton–Jacobi–Bellman equation over the whole action set. Piecewise linear flow benefits and penalties imply that a maximizing action can be chosen either at full exposure or at the safe boundary. For $\theta = 5$, $K_\theta = 0.042425$. Starting at 2 percent headroom, the institution initially unhedges fully, reaches this cutoff in 1.121 years if no shock occurs, then drops exposure to approximately 42.4 percent. It subsequently expands exposure safely as headroom grows.

This is a different threshold from the viability threshold in Proposition 2. Proposition 6 sets $g = 0$ and permits stress shortfalls; a firm below K_θ takes risk by choice under the stated objective. Proposition 2 permits an operating deficit and forbids such shortfalls. Mixing the two concepts would confuse a preferred gamble with an inability to finance a safe transition.

The abrupt change in (21b) depends on the linear control technology and the kinked loss penalty. Convex adjustment costs or smooth risk penalties can produce interior exposure choices. The result establishes a possible nonmonotone path and an exact benchmark, rather than a universal description of insurer policy.

Financing speed determines whether high penalties can eliminate risk-taking altogether. When $r < q$, a sufficiently high penalty makes the safe policy optimal at every capital level. When $r \geq q$, every finite penalty leaves a positive, possibly extremely small, low-capital region in which full exposure is optimal. The marginal value of moving away from very low capital is then unbounded along the safe value function. This conclusion depends on the perpetual financing opportunity and first-event objective; a fixed finite transition horizon or a recurring-shock model may change it.

To search beyond them, I solve an infinite-horizon Bellman problem. Within a time interval, the menu includes every point in a uniform grid of constant exposure actions, plus the exact evolving safe path for that interval. The adverse event’s timing within the interval is integrated analytically. Capital is interpolated linearly on a grid. Because capital is nondecreasing, the discretized Bellman equation can be solved from high capital down to low capital; self-interpolation terms are solved algebraically. Appendix B gives the recursion.

Shortfall penalty θ	Initial optimal unhedged share	Verified exact value	Numerical Bellman value
0	1.000	0.192308	0.192308
1	1.000	0.178755	0.178755
5	1.000	0.131097	0.131097
20	0.200	0.121590	0.121589
100	0.200	0.121590	0.121589

Table 2. Values are fractions of eligible exposure under criterion (19), not dollar welfare or estimated returns. The numerical action menu uses a 1/16-year interval, 3,201 capital points, and 801 constant exposure actions, augmented by the exact safe mode. Differences of less than 10^{-6} from the verified solution reflect discretization.

For $\theta = 5$, immediate full exposure is initially optimal, but remaining fully unhedged is not the best whole path. Figure 9 shows that exposure is subsequently reduced in a range of intermediate capital states. The option to change exposure produces a value above that of any of the four fixed recognition rules. At large penalties, the solution from the illustrative initial capital begins at the safe boundary. This is the substantive counterexample to recommending a universal exposure speed from the financing equations alone.




Figure 9. The left panel plots the verified continuous policy from Proposition 6 as a function of current capital, including states not reached from every initial condition. The dashed line is the strict safe boundary. The right panel compares the exact value on a dense penalty grid with the best of four continuous recognition rules; independent Bellman values are shown as open markers. The tiny risky region at the largest penalty may be visually compressed near zero.

Joint refinement of the time, capital, and exposure grids gives values of 0.131042, 0.131071, 0.131097, and 0.131097 at $\theta = 5$ for time intervals of 1/2, 1/4, 1/8, and 1/16 year. The final change is 8.35×10^{-8} ; the finest value is 1.47×10^{-7} below the exact optimum. The selected-policy Bellman residual is checked

below 10^{-11} , and the verified value's continuous HJB residual below 10^{-10} on a separate capital grid. Additional tests span $r/q = 0.5, 1, 1.5$ and four scaled penalties, including zero-cutoff and equality cases. These checks materially reduce numerical error concerns. They do not remove sensitivity to the assumed loss objective or first-event horizon.




Figure 10. The capital cutoff is expressed as a fraction of the full stress target. A cutoff of zero means the safe policy is optimal for every capital level; a cutoff of one means full exposure is optimal throughout. The dashed curve is the no-risk-taking boundary $\theta a/c = q/(q-r)$ for $r < q$. The map varies dimensionless ratios rather than asserting estimated parameter ranges for Taiwan.

14.6 Experiment 5 uncertainty about the savings that finance the buffer

The constant- c benchmark can be optimistic if savings fall when the currency shock becomes more likely. I therefore hold the aggressive policy fixed and let its annual pre-event saving switch between $c_L = 0.005$ and $c_H = 0.035$. The low state switches to high at intensity 1 per year; high switches to low at intensity 0.5. The stationary probabilities are $1/3$ and $2/3$, so mean savings remain 0.025. Retained capital grows at $v_s = \rho c_s$ in state s .

Three hazard assignments have the same stationary mean intensity of 0.10. The independent case uses $(\lambda_L, \lambda_H) = (0.10, 0.10)$. The adverse-association case uses $(0.22, 0.04)$; the opposite case uses $(0.02, 0.14)$. Only the assignment of the shock intensity to the cash-saving states changes. The process starts from its stationary state distribution.

Let $F_s(k)$ be expected discounted first-event shortfall starting in state s . For $k < a$, it satisfies the coupled backward equations

$$v_s F'_s(k) - (\delta + \lambda_s + \mu_s) F_s(k) + \mu_s F_{-s}(k) + \lambda_s(a - k) = 0, \quad F_s(a) = 0, \quad (22)$$

where $\mu_L = 1$ and $\mu_H = 0.5$. Appendix B derives this system. Since $v_s > 0$, the boundary at a is sufficient: once $k \geq a$, first-event shortfall is zero under the specified stress.

Relation of shock intensity to savings	Exact expected discounted shortfall in percent	Weekly simulation mean in percent	Monte Carlo standard error in percentage points
Constant savings and intensity benchmark	1.3553	1.3539	0.00331
State-dependent savings and independent intensity	1.4449	1.4509	0.00634
Shocks more likely when savings are low	1.5001	1.5090	0.00661
Shocks more likely when savings are high	1.3479	1.3570	0.00605

Table 3. Exact values solve (22) or, for the constant benchmark, (18). The Markov simulations use 150,000 paths, a weekly interval, and seed 271828. Their standard errors describe sampling uncertainty; a time-discretization allowance is checked separately. The constant benchmark uses the separate 500,000-draw experiment.

Adverse association increases the exact mean loss by approximately 3.8 percent relative to the stochastic-savings independent-hazard case, and by 10.7 percent relative to the constant-savings case. Reversing the association brings the mean loss slightly below the constant benchmark. This is an additional reason not to apply a generic claim that uncertainty must worsen the policy. The direction and timing of dependence matter.

The stationary hazard averages are equal, but the distribution of the first-event time need not be equal. Survival selects states differently. That selection is part of the mechanism represented in (22), not an accidental failure to hold a Poisson event distribution fixed. An experiment that required identical full marginal event-time distributions would answer a different question about copulas and would need a different construction.

The finite-horizon weekly simulation runs for 30 years. Since even the low state reaches the stress target within 20 years under the aggressive policy, no positive shortfall is truncated after 30 years. The coupled ordinary differential equations provide the primary numerical result; simulation is an independent check rather than the source of the reported exact means.




Figure 11. First-event expected shortfall across starting capital, computed from the coupled backward equations. The three state-varying cases have equal stationary average savings and event intensity. The constant benchmark is distinct. These curves are model expectations, not empirical confidence bands; the Monte Carlo cross-checks remain in Table 3.

15 Experiments on the interpretation of optimality

15.1 Experiment 6 separate the reward from the financing technology

Objective (19) rewards the full avoided expense flow cx , while only ρcx accumulates as capital. Those are different economic objects. To test the interpretation, replace the flow reward with bx , holding $\dot{k} = \rho cx$ unchanged. For example, $b = (1 - \rho)c$ rewards immediately distributable cash. This is an alternative research criterion, not a complete shareholder or household objective: retained capital has value and the first-event boundary still omits its continuation.

For an interior cutoff, the same HJB calculation gives

$$K(b) = a \left[\frac{q + (r - q)\theta\lambda a/b}{r} \right]^{r/(q-r)}, \quad r = \rho c/a. \quad (23)$$

The flow reward b therefore belongs in the cutoff independently of the resource flow ρc . With the common parameters and $\theta = 5$, the original 2 percent starting headroom yields the following results.

Flow reward share	Capital cutoff	Initial unhedged share
100%	4.242%	100%
80%	2.995%	100%
50%	1.286%	20%
20%	0.171%	20%

Table 4. Capital accumulation is identical across rows. The reward changes; the resources do not. The alternative objectives are not estimates of stakeholder utility.

Valuing only the 20 percent distributable share changes initial unhedging from 100 to 20 percent, and the cutoff from approximately 4.24 to 0.17 percent. A dense reward-by-capital grid makes the discontinuity visible rather than joining four selected points. The result identifies a policy-interpretation problem, not a mathematical error in Proposition 6.




Figure 12. Dense reward-by-capital sensitivity under the first-event objective. The orange line marks the 2 percent illustrative starting headroom. The color is a control choice, not a probability or confidence level.

15.2 Experiment 7 horizon and terminal continuation

A separate backward Bellman calculation replaces the infinite horizon with a finite decision interval. It retains the original event timing, savings and shortfall penalty. The first review used zero terminal continuation. A broader calculation also considers continuation at the value of following the safe policy after the horizon. The latter requires no favorable future risk-taking option, but it recognizes that the institution and its financing opportunity continue.

Horizon in years	Zero terminal value	Safe continuation value
0.25	20%	100%
1	20%	100%
3	20%	100%
5	100%	100%
10	100%	100%
20	100%	100%

Table 5. Initial unhedged shares at 2 percent headroom, with the same capital and loss parameters. Safe continuation means the exact remaining value of following the strict safe policy from the terminal capital stock; it is not a terminal payment equal to book equity.

The zero-terminal calculation chooses the safe 20 percent initial exposure at one and three years, while safe-policy continuation chooses full exposure even at a quarter-year horizon. Thus the earlier reversal is substantially a continuation-boundary result. It does not show that a short planning interval, by itself, makes gradual safety preferable. This opposing test is essential: a robustness critique should itself survive a plausible variation in the assumptions that generated it.

The new grid evaluates 80 quarter-year horizons, 801 capital nodes and 101 constant exposure choices, plus the exact adaptive safe branch. Both continuation choices are saved on a common capital-horizon surface. The original five-horizon and 50-year convergence results are retained. The verification theorem remains the primary result for the infinite-horizon problem; the finite calculations assess economically different boundaries.




Figure 13. Finite-horizon control surfaces under zero and safe-policy continuation. Both panels use identical capital dynamics and loss preferences; the orange line marks 2 percent starting headroom. The distinction is terminal continuation, not new empirical information.

16 Public balance sheets and market evidence

16.1 Data scope and the questions each source can answer

The empirical expansion combines 56,853 issuer-days for 14 selected financial holding tickers, with 56,555 usable issuer returns after 274 exchange-reference discontinuity flags. The equity history runs from 2010-01-04 to 2026-10-22. The primary files comprise monthly TWSE endpoint responses, including explicit no-data months. CBC daily FX supplies 4,673 observations from 2008. Twenty annual US Treasury CSV files supply two- and ten-year par yields. The official insurance balance sheet contains 472 month ends from May 1987 to August 2026. These counts describe different observation units and must not be added to claim an empirical sample size.

Fubon 2881 and Cathay 2882 remain the fixed life-oriented parent proxies. The original comparison comprises E.SUN 2884, Mega 2886 and First 2892. The broader comparison adds Hua Nan 2880 and SinoPac 2890. Other tickers enter individual-parent diagnostics. These are conglomerates with changing business mixes, not pure insurers and untreated bank controls. The chosen universe is not every Taiwanese insurer or every historical financial listing. Shin Kong 2888 is retained through its observed history; no cash payoff or merger total return is invented after observations end.

The longer history addresses regime coverage. Additional parents address comparison sensitivity. Neither supplies an insurer’s economic headroom, liability stress valuation or marginal hedge settlement. Repeating one daily currency regressor across 14 issuer rows does not create 14 independent currency shocks. The main inference is therefore on a time-series basket spread, with the number of dates reported.

16.2 Experiment 8 the long insurance balance sheet

The balance-sheet file permits a real historical perspective before the market episode. It records reported foreign assets, total assets, book equity and liability categories, with missing entries preserved. Reported asset components reconcile exactly in the downloaded vintage when all components are present; observed liability components also reconcile. The 2026 contract-liability entries are missing in the CSV despite the table’s new definition. An accounting residual is saved separately for audit and never substituted for an observed liability series.

Population and measurement changes matter. Foreign insurers’ Taiwan branches enter from 1989; postal life-insurance data enter from 1994. Portfolio-investment measurement changes to fair value from January 2006. IFRS 17 changes the insurance-liability definition in January 2026. Figure 14 starts in 1994 for visual comparability and marks the latest break, while the full earlier source history remains available. A level or growth break around these dates is not automatically an economic treatment effect (CBC 2026c).




Figure 14. CBC reported balance-sheet stocks and ratios. The full source covers 472 months; the plotted series starts in 1994 after postal-insurance inclusion. The dotted line marks January 2026. Book equity is not economic headroom; foreign assets are gross stocks.

Reported foreign assets fell by NT\$1.295 trillion in May 2025. An all-USD, fixed-quantity translation benchmark is NT\$1.419 trillion over that month; the remaining stock change is NT\$+0.124 trillion. Book equity changed by NT\$-0.183 trillion. The four amounts are not competing measures of the same loss.

Formally, let F_t be the reported foreign-asset stock in Taiwan dollars and S_t the end-month TWD-per-USD rate. The accounting decomposition is

$$F_t - F_{t-1} = F_{t-1} \left(\frac{S_t}{S_{t-1}} - 1 \right) + R_t. \quad (24)$$

The first term assumes an all-dollar foreign portfolio with unchanged foreign-currency quantity. The residual R_t combines other currencies, transactions, asset-price movements and reporting. It is not an identified net investment flow, intervention flow or hedge settlement. Likewise, gross foreign assets omit the offsets from policies, derivatives, reserves and other assets and liabilities. The decomposition

explains why a gross currency translation number cannot be relabeled an equity loss or a model capital shortfall.




Figure 15. Reported equity through August 2026 and the 2025 decomposition in equation (24). The benchmark assumes all foreign assets are dollars with unchanged foreign quantity. The residual is not a capital-flow estimate. The January 2026 equity movement is not identified as a reform effect.

A supplementary monthly regression relates the change in book equity, scaled by lagged total assets, to that gross translation benchmark on the same scale. It then adds changes in two- and ten-year Treasury yields and the TAIEX monthly log return on a common pre-2026 sample.

Monthly specification	Months	Coefficient	95% HAC interval
All available pre-2026 months	215	-0.171	[-0.320, -0.022]
currency only	191	-0.155	[-0.307, -0.002]
rates and market	191	0.095	[0.027, 0.163]

Table 6. Descriptive monthly accounting associations with six-lag HAC intervals. The regressor is a gross, all-dollar translation benchmark, not an identified currency surprise. Common-sample comparisons distinguish adding controls from changing the sample.

On the common sample, the currency-only coefficient is negative, while adding US yield changes and the equity market changes its sign to positive. This reversal makes the omitted-variable problem

observable. A currency movement can accompany bond-price and equity changes that push book equity in the opposite direction. The adjusted coefficient still cannot be interpreted as a hedge ratio, because the benchmark ignores currency composition, derivatives, liability sensitivities and reporting choices.

Even the sign of this accounting association must not be read as the economic FX sensitivity of an unhedged insurer. Currency moves co-occur with bond-price and equity changes; book accounting and liability values also matter. The aggregate series cannot recover the counterfactual cash cost c or economic headroom k . A superficially precise aggregate coefficient can therefore answer the wrong structural question.

16.3 Experiment 9 longer and broader equity exposure diagnostics

For date t , define $A_t = -100\Delta \log S_t$, so positive values indicate Taiwan-dollar appreciation. The dependent variable Y_t is the equally weighted clean price return of the two life-oriented parents minus that of the comparison parents. The descriptive specification is

$$Y_t = \beta_0 + \beta_M M_t + \beta_A A_t + \varepsilon_t, \quad (25)$$

where M_t is the TAIEX price return in percentage points. A rate-control version adds changes in US two- and ten-year par yields. US quote dates strictly precede the Taiwan equity date; changes in the available levels are taken across Taiwan exchange dates so a long holiday does not discard intervening US moves. A same-day Taipei FX close is later than the regular 13:30 equity close. Consequently, equation (25) is an association across overlapping information windows, not an ex ante trading signal or causal structural currency shock.

The main stock-return measure excludes issuer-days where the previous close differs by more than NT\$0.051 from the reference price implied by the exchange's reported price change. This retains the earlier review's cleaning rule. The alternative reference-return measure retains the exchange's new price basis on discontinuity dates. Neither is a complete, reinvested total-return index. Flags and exclusions remain inspectable, and zero returns are not substituted for missing observations.

The original 2020-onward five-parent test remains a separate comparison: 1,612 complete dates give a coefficient of -0.210, with a five-lag HAC interval of [-0.556, 0.135]. Removing the known May stress window gives -0.072 and [-0.341, 0.197]. These earlier estimates remain in the replication unchanged; the larger sample, broader basket and ten-lag uncertainty convention below are explicitly different specifications.

Sample or specification	Dates	FX coefficient	95% HAC interval
2010 onward	3992	-0.274	[-0.471, -0.077]
2010–2019	2390	-0.329	[-0.575, -0.083]
2020–2024	1180	-0.022	[-0.337, 0.294]
2025 onward	422	-0.337	[-0.874, 0.199]
excluding May stress	3983	-0.221	[-0.399, -0.043]
reference returns	4116	-0.289	[-0.472, -0.106]
prior-US yield controls	3990	-0.256	[-0.468, -0.045]

Table 7. Broader five-parent comparison; the original three-parent results remain in the replication and appendix. Daily intervals are normal-approximation, ten-lag Newey–West intervals with a finite-sample factor. All estimates are retrospective.

The longer history strengthens evidence for a negative daily conditional association. The full-period broader-basket coefficient is approximately -0.274, with an interval of [-0.471, -0.077]. It remains negative after excluding the May stress window (-0.221, interval [-0.399, -0.043]) and after adding prior-US yield controls (-0.256, interval [-0.468, -0.045]). The earlier 2020-onward review therefore understated what the longer historical daily data can establish: the pooled association is not solely the selected May episode.

That stronger result is still regime dependent. The 2010–2019 coefficient is -0.329, while the 2020–2024 coefficient is close to zero and imprecise; the 2025-onward interval is also broad. Weekly aggregation gives -0.072 across 849 complete weeks, with an interval of [-0.397, 0.253]. Thus the daily full-period result does not establish one invariant exposure coefficient across periods and observation frequencies. Timing, market anticipation, business-mix changes and other common news remain plausible contributors.




Figure 16. Broader comparison coefficients and approximate 95 percent HAC intervals. The coefficient is relative-return percentage points per 1 percent log appreciation. Sample and control choices are shown rather than selected for significance.

Figure 17 shows the full sequence of trailing three-calendar-year estimates on quarterly evaluation dates. The windows overlap. Figure 18 gives every selected parent's own market-controlled currency coefficient using its available history. These disaggregations test whether an aggregate coefficient conceals regime and business-mix differences. They do not supply a constant historical insurance classification for acquisitions, mergers or renamed groups.




Figure 17. Trailing three-calendar-year currency coefficients with approximate 95 percent ten-lag HAC bands. Overlapping quarterly evaluation windows are descriptive, not independent observations.




Figure 18. Market-controlled coefficients and 95 percent HAC intervals for the 14 selected financial parents. Orange marks the two primary life-oriented proxies. Histories differ by listing and delisting; names do not imply unchanged business mixes.

Additional checks use separate appreciation and depreciation slopes, weekly aggregation of exchange-reference returns, leads and lags of FX, leave-one-calendar-year-out estimates and paired moving-block resampling of the broader basket regression. The bootstrap uses 2,000 resamples at each of three block lengths. Its percentile intervals describe stability under resampling the observed time-series pairs; they are not confidence intervals for a reform effect.

Block observations	Resamples	2.5th percentile	97.5th percentile
5	2000	-0.464	-0.066
20	2000	-0.467	-0.082
60	2000	-0.482	-0.071

Table 8. Percentile intervals for the daily appreciation coefficient under historical moving-block resampling. Complete rows are resampled in contiguous blocks; gaps created by the cleaning exclusions mean contiguity refers to observed complete rows.

These checks should be read jointly, rather than selecting the period, basket or uncertainty estimator with the strongest result. A larger panel can make a pooled coefficient more precise while leaving the model's structural mapping unidentified. The observed slope is conditional equity-price exposure, not the stress-capital coefficient α , an insurer hedge ratio, or a causal influence on the CBC.

16.4 Experiment 10 the May episode and alternative event windows

The original five-parent May 2 and May 5, 2025 comparison remains informative: the life-oriented basket's summed clean daily price returns are -5.76 percentage points, versus -0.64 for the original bank-oriented basket. A 2020–2024 market-only fit gives a cumulative abnormal spread of -5.58 points. Its retrospective overlapping-window lower-tail rank is 0.32 percent. That rank is not an independent-test p-value: the episode was known before the exercise and the placebo windows overlap.

Figure 19 uses the newly broadened comparison but preserves the observed episode. The cumulative paths compound daily basket returns, whereas the original event statistics sum daily returns and market-model residuals. These definitions differ deliberately and are not silently interchanged. Currency news cannot be isolated from tariff news, interest rates, portfolio flows and changing policy expectations in that short window.




Figure 19. The known May 2025 stress episode, with the broader comparison. The upper panel compounds clean price returns; the lower panel shows CBC FX closes. Equity and FX closes occur at different times.

The December 23–24, 2025 proposal window sits near the middle of the original placebo distribution. That does not establish either no policy effect or an improvement in economic solvency. Two planned windows remain excluded. July 1, 2025 has a Fubon reference-price reset; the raw close declines from 87.3 to 83.2 while the exchange reports zero change against the new reference. The February 12–13, 2026 window has no exchange observations in the retrieved index series; extending to the post-holiday

reopening would mix additional news. Missing event evidence is retained as a result, not filled with zero returns.

16.5 Experiment 11 genuinely prior-information forecast evaluation

The first review’s 2025-onward diagnostic conditioned on the contemporaneous market return. Adding lagged FX slightly worsened RMSE, from 1.3483 to 1.3591 return percentage points. Since the benchmark already uses the realized market return, it is best described as a conditional fit diagnostic, not an equity-open forecast.

The expanded evaluation instead refits once at the beginning of each year from 2015 onward, using only preceding years. The baseline includes the preceding exchange-date basket spread and market return. The augmented model adds the last CBC daily FX move observed before the current equity date. Predictions are evaluated on the same complete rows. Training endpoints are checked to precede every evaluation date. The specification is fixed across years; there is no search over trading rules.

Across 2,720 evaluated dates, baseline RMSE is 1.0155 return percentage points and augmented RMSE is 1.0164. The augmented model improves RMSE in 6 of 12 annual evaluations. The mean augmented-minus-baseline squared-error difference is 0.00187, with a 20-lag HAC interval of $[-0.00111, 0.00484]$. These are retrospective chronological results, not an untouched test: the May episode and earlier findings were already known when this extension was specified.




Figure 20. Annual RMSE difference from fixed expanding-window specifications using only previous information. Negative bars favor adding lagged FX; positive bars favor the baseline. The partial 2026 evaluation ends at the observed source cutoff.

No portfolio return, transaction-cost-adjusted alpha or trading recommendation follows. Predictability of parent equity returns is also not a necessary implication of the transition-financing theory. A weak forecast result can coexist with meaningful balance-sheet exposure and largely efficient prices.

17 Historical account replays and conditional forward experiments

17.1 Experiment 12 transport the rule into continuing marked currency movements

The first-event model deliberately stops when an adverse event occurs. A marked-account replay changes that economic environment. Let $Q_t = S_t/S_0$ measure the domestic price of a fixed foreign-currency block relative to its initial price. Exposure x_t is chosen using prior information, and the proportional turnover charge is κ . The continuing account follows

$$K_t = K_{t-1} - \kappa Q_{t-1} |x_t - x_{t-1}| + (g + \rho c x_t Q_{t-1}) \Delta t + x_t (Q_t - Q_{t-1}). \quad (26)$$

The common assumptions are $K_0 = 0.02$, $a = 0.10$, $c = 0.025$, $\rho = 0.8$, $g = 0$, and a 25-basis-point turnover charge. The safe rule finances the current adjustment fee before applying its prescribed stress constraint. Positive currency marks are fully retained. A negative account ends that research path and its value is frozen; that boundary is not an actual insurer default, liquidation recovery or continuing insurance profit-and-loss account.

The cash saving, capital and adjustment charge are frozen illustrative parameters. Actual bond returns, liability values, collateral calls, hedge forward points, maturities and funding are omitted. Thus the replay tests whether the original rule transfers to this specified account, not whether it is an optimal real hedge. It is not an investment-performance backtest.

Rule	Five-year windows	Account crossed zero	Median terminal account
hedged	166	0.00%	2.00%
safe	166	0.00%	2.82%
transported optimal	166	80.12%	-0.08%
unhedged	166	66.87%	-0.06%

Table 9. The 166 monthly-start five-year histories overlap. Fractions crossing zero are properties of these hypothetical accounts on these paths, not independent estimates of an insurer-failure probability.

The transported first-event policy crosses zero more often than permanent full exposure under this specification. After repeated losses, it can return to a low-capital full-risk region; gains that reach the cutoff induce it to cut exposure. This asymmetry explains how a rule can be optimal for a terminating event problem and perform poorly when losses recur. The exercise challenges portability, not the proof of Proposition 6.

The daily stress-safe account has no zero crossing in the observed five-year histories. That is partly built in: a 10 percent one-step stress allowance exceeds the observed single-day losses in this source history. The result does not validate the chosen stress, establish the absence of liquidity risk, or certify an insurer's solvency.




Figure 21. The left panel uses every overlapping five-year history and plots terminal-account distributions. The right panel replays the actual 2025 FX path. Negative accounts freeze at their first crossing; the hedged account is a relative baseline that omits the wider insurance business.

17.2 Experiment 13 adjustment frequency and proportional cost

The expanded sensitivity evaluates stress losses of 2, 5 and 10 percent, adjustment every 1, 5 or 20 observations, and turnover charges of 0, 25 and 50 basis points. It uses 190 overlapping three-year windows. The parameters form a complete 27-case comparison rather than a selection of favorable policies. Lower prescribed stresses permit larger initial exposures, so they can become more fragile when several daily moves arrive before adjustment.

At the 5 percent stress and 25-basis-point charge, daily adjustment has no observed crossings; adjustment every 5 and 20 observations gives approximately 4.74 and 6.32 percent. The 10 percent cases remain below the crossing boundary on these historical windows. The 2 percent cases are substantially more fragile. This is an interaction among the chosen stress, cost and adjustment opportunity, not evidence that the least frequently traded real hedge necessarily fails.




Figure 22. All 27 stress, frequency and charge combinations for the stress-safe rule on 190 overlapping three-year windows. Cell values are conditional historical crossing fractions, not estimated insurer-default probabilities.

17.3 Experiment 14 forward block scenarios

Forward conditional scenarios resample observed daily log currency changes in blocks of 1, 20 or 60 observations. Each calibration-pool and block-length pair uses 6,000 common currency paths across policies, 756 observations and a fixed 252-observation model year. The two pools end before 2025 or at the observed cutoff. Seed 20261026 fixes the calculation. These are alternative conditional environments, not probabilistic forecasts of a future insurer balance sheet.

Pool	Block	Unhedged account	Transported rule	Stress-safe
pre 2025	1	52.57%	60.47%	0.00%
pre 2025	20	61.30%	68.52%	0.00%
pre 2025	60	62.52%	69.55%	0.00%
through cutoff	1	56.20%	63.28%	0.00%
through cutoff	20	63.42%	69.93%	0.00%
through cutoff	60	63.57%	70.67%	0.00%

Table 10. Fractions of simulated research accounts crossing zero. Rows differ in their historical support and preserved dependence. Six thousand scenarios do not constitute six thousand independent historical currency regimes.

Longer blocks change the comparison relative to independent daily resampling. Including the recent stress history also matters. A stationary historical bootstrap cannot invent a daily shock absent from its pool, estimate a new policy regime or fully represent rare tails. Consequently, the stress-safe rule's zero in this experiment is primarily a property of its one-step design and the support of the resampled moves.




Figure 23. Block and pool sensitivity, with conditional terminal-account percentiles. The right-panel intervals are 5th–95th scenario percentiles, not empirical confidence intervals. The model omits actual insurer assets, liabilities and hedge quotations.

17.4 Experiment 15 stress absorption and the cost of the next action

Named paths address the missing-tail problem directly: an immediate 10 or 12 percent loss, two successive 3 percent losses, a gradual 12 percent loss over a year, and a gradual 10 percent gain. The same account and cost conventions apply. These paths have no estimated likelihood.

Specified path	Stress-safe	With exit buffer	Unhedged	Transported
One-day 10% loss	Yes	No	Yes	Yes
One-day 12% loss	Yes	Yes	Yes	Yes
Two successive 3% losses	No	No	Yes	Yes
Gradual 12% loss	No	No	Yes	Yes
Gradual 10% gain	No	No	No	No

Table 11. Yes means the hypothetical account crosses zero. A path within the nominal 10 percent stress can still reveal missing exit liquidity; a 12 percent jump is outside that stress.

The ordinary stress-safe account absorbs a one-day 10 percent loss but can cross zero when it subsequently pays to close exposure. The capital needed for the loss and the capital needed for the next action are separate. If K denotes capital after paying the current entry-adjustment charge, a sufficient condition to absorb a loss a and close the remaining exposure is

$$K \geq \{a + \kappa(1 - a)\}xQ. \quad (27)$$

The first term finances the currency loss; the second finances closing the post-loss exposure. At $a = 0.10$ and $\kappa = 0.0025$, the effective requirement is 10.225 percent of current exposed value. The added exit-buffer rule survives the named 10 percent path, while still failing the 12 percent jump outside its stress. This is a derived correction for the proportional-cost account, not an estimated liquidity requirement for Taiwanese insurers. Fixed closeout fees, margin jumps or insufficient market depth would require a different bound.

18 From insurer capital to exchange rate adjustment

The core model conditions on the currency stress. To explain why it could matter for macroeconomic adjustment, consider a separate, reduced-form policy problem. Let X be aggregate newly unhedged exposure, k the corresponding aggregate headroom, and a_0 a currency adjustment preferred in the absence of a financial stability concern. The policymaker chooses $b \in [0, a_0]$ to minimize

$$\frac{\chi}{2}(b - a_0)^2 + \frac{\nu}{2}(bX - k)_+^2, \quad \chi > 0, \nu > 0. \quad (28)$$

The first term represents the cost of departing from the unconstrained adjustment. The second represents concern about a financial stress shortfall. When $a_0X \leq k$, the choice is $b^* = a_0$. Otherwise,

$$b^* = \frac{\chi a_0 + \nu X k}{\chi + \nu X^2}. \quad (29)$$

For $X > 0$ in the shortfall region, more headroom permits a larger adjustment. More unhedged exposure reduces the chosen adjustment. The first derivative is $\partial b^* / \partial k = \nu X / (\chi + \nu X^2) > 0$. The sign of the second derivative follows from $k < a_0X$; Appendix A supplies the calculation. A capital path satisfying $k \geq a_0X$ removes this particular source of distortion entirely.

Equation (29) is a mechanism, not evidence about the CBC's actual decisions. The policymaker may be unable to choose the exchange rate directly. Export prices, imported inputs, inflation, foreign demand, cross-border flows, and geopolitical risk also enter the adjustment problem. Nor does a high current-account surplus itself identify a_0 . Calling (28) a measured social welfare function would assign unsupported empirical content to χ and ν .

The extension nevertheless clarifies a conditional causal sequence. Lower recognition affects exposure only if a behavioral or regulatory constraint responds to it. Greater exposure raises incremental losses only to the extent that liabilities and other assets do not offset them. Those losses affect exchange rate policy only if the authority assigns them material weight and has an effective instrument. Each link is contestable and separately testable. Observing a hedge reduction and an exchange rate movement establishes none of them by itself.

19 Empirical implications and a design that could reject the mechanism

The broad candidate outcomes in this earlier design are superseded for the primary institutional test by the common-equity cash-issuance specification above. They remain supporting possibilities, not additional primary significance tests.

19.1 The central measurement task

The first empirical output should be an institution-level reconciliation rather than a reform regression. For each insurer and quarter, reconstruct the eligible foreign exposure, derivatives by instrument and maturity, foreign currency policy liabilities, amortized-cost debt, unrecognized foreign exchange differences, relevant reserve accounts, dividends, equity issuance, and economic asset and liability sensitivities. Published annual and interim financial statements and their foreign currency sensitivity notes are natural starting points. Some crucial inputs, including instrument-level hedge cash costs and stress-consistent liability revaluation, may require supervisory access. The public aggregate and equity evidence assembled above does not supply that institution-level reconciliation.

The model's k would require a consistent valuation bridge. Begin with a stated economic valuation of assets and liabilities, subtract a fixed required floor, and subtract stress losses associated with positions outside the newly unhedged block. Accounting recognition adjustments must be reversed consistently. Show the reconciliation to reported equity and eligible capital rather than silently treating any of them as equivalent.

Likewise, identify c from incremental net cash flows. Hedge expense cannot be read mechanically as an avoidable resource cost if it includes settlements that offset losses elsewhere. Derivative cash settlements, forward points, interest differentials, collateral funding, taxes, and associated foreign asset returns should be accounted for on a consistent counterfactual basis. The estimated ρ must describe actual incremental retention after distributions. Legal reserve allocation rates alone cannot identify it.

19.2 Predictions that distinguish the transition mechanism

The first prediction is a resource reconciliation. Conditional on an identified hedge reduction, actual cash saving, and fixed exposure, economic headroom should accumulate with retained net surplus rather than with the announced amount of an accounting reclassification. If reported profit rises while retained economic resources do not, the proposed financing channel has not been established.

The second prediction concerns the initial lower tail. Conditional on similar stress exposure and saving rates, institutions below a reconstructed transition target should be more likely to need issuance, reduce distributions, change portfolios, or miss the transition deadline. The relevant nonlinear variable is $(z_H - k_{i0})_+$, not merely a linear interaction with average capital. If accurately measured weak firms complete the same safe transition without additional resources, the model's restrictions or measurements are wrong.

The third prediction separates reported stabilization from economic stress. High eligible amortized-cost exposure can predict a larger change in reported exchange rate sensitivity even when the economic sensitivity is unchanged. A subsequent reduction in true sensitivity would require a change in liabilities,

exposure, or loss absorption. This distinction is observable in principle through parallel reported and common-stress balance sheets.

The fourth prediction is about timing. A change in economic loss absorption should be larger after a period of actual retention than immediately at the accounting effective date, unless fresh resources arrive at once. This prediction can be rejected by an immediate cash or liability revaluation channel. It should not be protected by redefining every reported reserve increase as retained cash.

19.3 An exposure design and its identification limits

A preliminary panel specification could use an exposure fixed before the December 2025 announcement:

$$Y_{it} = \eta_i + \gamma_t + \sum_{\ell \neq -1} \beta_\ell (E_{i,\text{pre}} \mathbf{1}\{t - T = \ell\}) + \Gamma' W_{it} + \varepsilon_{it}. \quad (30)$$

Here $E_{i,\text{pre}}$ is the pre-announcement share of debt eligible for the recognition treatment, Y is a precisely defined reporting or behavioral outcome, and W contains predetermined or clearly justified controls. The announcement date and accounting effective date would be coded separately. December and February policy components should not be collapsed into a single event without a reason.

The coefficients in (30) describe differential changes by prior exposure. A causal interpretation requires that, absent the reform, high- and low-exposure insurers would have followed comparable outcome paths after accounting for relevant shocks. That assumption is demanding. Eligibility reflects earlier portfolio and hedging choices; the 2025 appreciation can already have affected capital and hedges. IFRS 17 implementation and the new solvency framework coincide with the period. Liability duration, guaranteed returns, ownership, and access to group capital can produce differential responses unrelated to the accounting change.

A pre-announcement eligible share is therefore not automatically an instrument. Its exclusion restriction would fail if it directly predicts exchange rate losses, duration exposure, or the response to the other regulatory changes. The same caution applies to using an insurer's own hedge ratio as an explanatory variable: it is an endogenous choice that may respond to deteriorating capital or expected policy support.

The design should publish pre-trends, reconciliation quality, and common-support diagnostics before emphasizing a post-treatment coefficient. Use outcomes that distinguish cash retention, exposure, and reporting; inspect firms individually; report leave-one-firm-out estimates; and account for the small and concentrated insurer population. Foreign currency policy liabilities and instrument-specific transmission are useful heterogeneity checks, not unquestionable placebos, because the reform may affect them too.

There is no obvious untreated domestic insurance sector with the same liabilities and foreign portfolio. International comparisons would face different regulation, currencies, and dollar funding arrangements. A sector-level before-and-after comparison is descriptive. An interrupted time series cannot by itself resolve the simultaneous shocks.

19.4 Testing the policy link separately

The macro extension requires different evidence. One would need high-frequency currency surprises, documented intervention where available, instrument-specific hedge adjustments, and predetermined insurer sensitivities. A joint response of insurer valuations and the currency could establish exposure to common news; it would not identify an insurer-induced policy response. Policy communications and interviews may clarify the decision mechanism but cannot substitute for an identification argument.

The CBC's distinction between offshore forwards and onshore swaps suggests an explicit test of transmission rather than a single hedge-total variable. A decline in offshore non-deliverable forward exposure need not create the same onshore dollar demand as settlement of a swap. If the claimed exchange rate effect appears equally across instruments despite their different mechanisms, that would weaken the proposed portfolio-flow channel.

The empirical paper should therefore make two contributions in sequence. First establish whether retained cash actually relaxes institution-level stress constraints. Then investigate whether those constraints materially influence currency adjustment. Combining both into one aggregate narrative before testing either would obscure the most informative possible rejection.

20 Criticisms and alternative mechanisms

20.1 The prescribed stress may be the wrong risk criterion

The safe constraint assigns decisive importance to one loss size at every date. A policymaker might rationally accept small expected losses in exchange for much lower ongoing hedge costs. Experiment 4 demonstrates that conclusion within the paper's own framework. Conversely, a 10 percent stress may understate a relevant tail event. The model supplies a frontier conditional on a ; it cannot select the socially correct stress or loss penalty.

Repeated shocks are an additional concern. Once an adverse event uses capital, an institution must rebuild from a different state. A strategy that safely finances one transition may not be viable in an environment with recurring jumps, persistent exchange rate drift, or changing liability valuations. The historical replays and forward scenarios above make the portability problem concrete. The exact recurring-state control experiment earlier in this revision supplies a new optimum under its own explicit assumptions; a full asset-liability model would still be required before making long-run solvency claims.

20.2 Hedging itself can create liquidity risk

Conventional hedges can require collateral or repeated rollover. Eliminating part of that demand may reduce a liquidity shortfall in a dollar funding disruption. McGuire et al. (2021) document the relevance of such arrangements in emerging Asia. The present safe constraint only values the hedge's reduction of currency losses; it gives the hedge no funding or collateral risk. Its policy comparison can therefore favor retaining hedges too strongly in a liquidity crisis.

A joint extension would require a liquidity state, collateral rules, instrument maturities, and a funding shock. Its objective would price both currency losses from unhedging and funding losses from hedging. Those costs should not be compressed into the same a without specifying the joint event. The direction of an optimal hedge change can reverse across appreciation and dollar-liquidity stresses.

20.3 Liabilities may absorb the currency movement

Foreign currency policies offer a natural offset for the insurer, but transfer currency risk to policyholders whose consumption is in domestic currency. That transfer may be an informed choice rather than a distortion. A welfare analysis would need household portfolios, consumption baskets, policy guarantees, and distributional weights. Lower insurer exposure alone is insufficient evidence that national currency risk has fallen.

Domestic currency liabilities can also have market values that respond to interest rates or inflation. If the currency movement is accompanied by favorable liability revaluation, the incremental ax approximation overstates the loss. If it coincides with credit or equity losses, it understates it. The institutional application must stress both sides of the balance sheet jointly.

20.4 Accounting relief can preserve going concern value

A forced sale triggered by transient reported losses may destroy value even when ultimate contractual claims are supportable. In that case, changing recognition can prevent resource destruction and raise the true baseline surplus or initial headroom. The neutrality benchmark deliberately holds those channels fixed. Evidence of avoided fire sales or cheaper funding would reject its applicability, not contradict its algebra.

The positive findings reported by Wang (2026) under combined accounting and reserve mechanisms are compatible with such an effect. They are also compatible with sufficient actual retention compensating for greater unhedged exposure. To distinguish these explanations, a replication must keep cash-flow assumptions, distribution rules, and capital definitions consistent across policy cases. A favorable simulated accounting path alone would not identify which mechanism produced it.

20.5 General equilibrium and outside financing may dominate

The model takes hedge costs as given. If many insurers reduce hedges, market prices and dealer balance-sheet demand can change, eroding the marginal saving. A lower c slows the financing path and raises the critical capital level. If instruments become cheaper as demand falls, remaining hedges may also become more attractive. A sector model should allow both effects.

New capital, portfolio diversification, liability repricing, mergers, and a deeper domestic long-term bond market could all alter the result. Equation (10) is useful precisely because it states the missing resource when those channels are absent. It does not imply that an up-front public recapitalization is the best source. Private issuance can be costly; a public guarantee can subsidize risk; redirecting assets can change returns; repricing liabilities can shift losses to customers. The paper does not rank those alternatives without their costs.

20.6 Novelty and the standard for an empirical contribution

The core stock-flow distinction is elementary, and the proofs use standard methods. The paper would not justify a strong publication claim merely by placing a recent reform into this model. Its strongest research prospect is to show that the deadline requirement and capital distribution explain institution-level responses that existing accounting or risk-management models miss. That requires original balance-sheet reconstruction, credible variation, and specialist scrutiny of the institutional rules.

The computational work adds model-based evidence on optimization, adverse timing, continuation value and exit liquidity. The larger public panel improves regime and comparison coverage. Neither substitutes for that institution-level empirical contribution. In particular, simulated paths cannot establish the fraction of Taiwanese insurers below the threshold, the size of a real recapitalization budget, or the exchange rate Taiwan can safely sustain.

21 Conclusion

The transition funding gap is a conditional implementation requirement, not an assumption-free balance-sheet statistic. Its definition needs a destination, deadline, stress system, contractual cash calendar and admissible financing channels. The computed frontier separates cash from loss absorption and can distinguish institutions that have identical opening capital, exposure and liquidity summaries.

The expanded-access test removes the synthetic timing gap when the pre-existing claim can be monetized on sufficiently favorable terms. A remaining empirical timing mechanism therefore needs a demonstrated restriction on saleability, pledging or legal use, rather than low current cash alone. The derivative-disclosure audit likewise shows a liability maturity total falling while aggregate nominal volume rises; the former is not a hedge-quantity proxy.

The completed calculations also establish limits. A full future settlement ladder already contains the timing information. Forward carry is not automatically deadweight expense. Book accounting relief does not itself create cash or economic capital. An exact recurring-risk solution differs from an imported first-event rule, while the favorable net flow and resolution objective remain assumptions rather than estimated insurer economics.

Six primary-report bridges supply an auditable institutional starting point, and the earlier market and stress evidence remains intact. Instrument-level hedge cash flows, economic liability values and legal resource access are still missing. Actual institutional gaps and their incremental predictive performance are consequently unidentified. The proposed financing prediction, harder liquidity benchmarks and explicit rejection criteria state what would establish or weaken the mechanism. A liquidity bridge can finance a payment timing problem; it cannot, by itself, repair an economic loss.

Appendix A proofs and analytical details

Proof of Proposition 1

Define $h(k) = \min\{1, k/a\}$ for $k \geq 0$ and $f(k) = g + vh(k)$. The function f is Lipschitz on the nonnegative domain. For any admissible capital path k , the stress constraint implies $x_t \leq h(k_t)$ almost everywhere. Thus $\dot{k}_t \leq f(k_t)$. The scalar differential inequality comparison theorem, or its direct Gronwall proof applied to the positive part of $k_t - k_t^*$, yields $k_t \leq k_t^*$ until the end of their nonnegative domain.

Since h is increasing, $x_t \leq h(k_t) \leq h(k_t^*) = x_t^*$ almost everywhere. Any strictly positive weighted integral of cx_t is therefore weakly smaller than that of cx_t^* . Before k_t^* reaches a , its equation is $\dot{k}_t^* = g + rk_t^*$. Multiplying by e^{-rt} and integrating gives (5). At $k = a$, equation (4) changes to $\dot{k} = g + v$. If the maximal path hits zero with negative drift, no other admissible path can continue beyond that bound without another resource source. This proves the claimed ordering on viable horizons.

Proof of Proposition 2

For $-v < g < 0$, the critical point $k^{\text{crit}} = -g/r$ lies strictly between zero and a . Equation (5) becomes $k_t^* = k^{\text{crit}} + (k_0 - k^{\text{crit}})e^{rt}$. If $k_0 > k^{\text{crit}}$, its distance above the critical point increases, and solving $k_t^* = a$ yields (7). The post-target drift $g + v$ is positive, so full exposure remains safe.

At equality, the path is stationary. Below the critical point, the distance below it increases. Solving $k_t^* = 0$ gives (8) for $k_0 > 0$. At zero, the stress constraint forces $x = 0$ and the capital drift is $g < 0$, violating nonnegativity immediately. Proposition 1 implies that no smaller capital path can do better. If $k_0 = 0$, failure is immediate. The cases $g = 0$, $g > 0$, and $g + v \leq 0$ follow directly from the same linear equation and the drift at full exposure.

Proof of Proposition 3

For $-v < g \leq 0$ and an initial state $K < a$ above the critical point, the first time that the maximal safe path reaches a is at most H if and only if

$$k^{\text{crit}} + (K - k^{\text{crit}})e^{rH} \geq a.$$

Rearrangement gives $K \geq z_H$. Since z_H lies between the critical point and a , initial capital at or above a also satisfies this condition. An up-front injection J changes K to $k_0 + J$. Necessity follows from maximal-path dominance; sufficiency follows by using that path from $k_0 + [z_H - k_0]_+$. Equality attains the deadline exactly when a positive injection is needed. The monotonicity in H follows by differentiation.

Proof of Proposition 4

Each institution requires at least $[z_H - k_{i0}]_+$ by Proposition 3. Since resources cannot be transferred between institutions, those requirements must be financed separately, yielding (11). The function $\phi(k) = [z_H - k]_+$ is convex. Jensen's inequality gives $\sum_i w_i \phi(k_{i0}) \geq \phi(\bar{k}_0)$. The defining convex-order property of a mean-preserving spread gives the dispersion result.

With free transfers, if $\bar{k}_0 \geq z_H$, reallocate headroom so that every institution has at least z_H without adding resources. If $\bar{k}_0 < z_H$, an external amount $z_H - \bar{k}_0$ permits exactly that allocation. A smaller amount is impossible because the weighted total would fall below z_H . This proves attainability of the pooled bound.

Proof of Proposition 5

For $k < \alpha a$, rule (13) gives $ax_\alpha(k) = k/\alpha$, so the shortfall is $k(1/\alpha - 1)$. For $\alpha a \leq k < a$, exposure is one and shortfall is $a - k$. For $k \geq a$, shortfall is zero. The first branch is increasing and the second decreasing, and both equal $a(1 - \alpha)$ at $k = \alpha a$. This proves (15) and the maximum statement.

For $g = 0$ and $k_0 < \alpha a$, capital before full exposure satisfies $\dot{k} = vk/(\alpha a)$, giving $k_t = k_0 e^{vt/(\alpha a)}$. Setting $k = \alpha a$ gives (16). Subsequently $\dot{k} = v$ until reaching a , giving the additional time in the text. For $\alpha = 0$, exposure is one immediately and $k_t = k_0 + vt$.

Derivation of the two rule loss comparison

Under immediate full exposure, shortfall at time $t < T_A$ is $D - vt$ and is zero thereafter. Its expected discounted value is $\int_0^{T_A} \lambda e^{-qt}(D - vt)dt$. Integrating the constant and linear terms gives (18).

Under the safe path with $g = 0$, exposure is $(k_0/a)e^{rt}$ up to τ_S and one afterward. Integrating $cx_t e^{-qt}$ over those intervals gives (20). The shortfall penalty is zero for that path. Subtracting its value from $B_A - \theta L_A$, and using $L_A > 0$ for $0 < k_0 < a$ and $\lambda > 0$, yields (21). No comparison with every admissible loss-tolerant policy follows from this two-rule calculation.

Proof of Proposition 6

For $g = 0$, the continuous HJB equation for (19) on $0 \leq k < a$ is

$$qV(k) = \max_{0 \leq x \leq 1} \{cx - \theta\lambda(ax - k)_+ + vxV'(k)\}, \quad V(a) = c/q. \quad (31)$$

Increasing initial capital weakly reduces losses under the same exposure path, so the value is nondecreasing. The candidate below has a nonnegative derivative. For any such candidate, the expression being maximized is linear and increasing on $0 \leq x \leq k/a$. On $k/a \leq x \leq 1$, its slope is $C(k) = c + vV'(k) - \theta\lambda a$. Thus a maximizing choice can be made at either k/a or one.

When $\theta \leq c/(\lambda a)$, consider the full-exposure value $V_A(k) = c/q - \theta L_A(k)$, with (18) evaluated from k . Direct differentiation gives $V'_A(k) = \theta\lambda(1 - e^{-q(a-k)/v})/q \geq 0$. Hence $C(k) \geq c - \theta\lambda a \geq 0$, so full exposure maximizes the HJB expression. This bounded value satisfies (31).

Now suppose $\theta > c/(\lambda a)$. Write $y = k/a$. For $r \neq q$, the value of following the safe path from capital $k > 0$ is

$$V_S(k) = c \left[\frac{r}{q(r-q)} y^{q/r} - \frac{y}{r-q} \right], \quad V'_S(k) = \frac{c}{a(r-q)} (y^{q/r-1} - 1). \quad (32)$$

For $r = q$, the continuous limit is $V_S(k) = (c/r)y(1 - \log y)$ and $V'_S(k) = -(c/(ar)) \log y$. In all cases $V_S(0) = 0$ by continuity. The slope $C_S(k) = c + vV'_S(k) - \theta\lambda a$ decreases strictly with k and is negative

at a . When $r \geq q$, it tends to positive infinity as k tends to zero. Its unique zero is the positive cutoff stated in the proposition.

When $r < q$, its limit at zero is $cq/(q-r) - \theta\lambda a$. If this is nonpositive, $C_S(k) \leq 0$ at every state, the safe action maximizes (31) everywhere, and $K_\theta = 0$. If the limit is positive, the unique interior zero is (21a). Thus, whenever a positive interior cutoff exists, the safe action maximizes (31) for $k \geq K_\theta$. At that boundary, the safe HJB equation implies $qV_S(K_\theta) = \theta\lambda K_\theta$.

For $k < K_\theta$, define the value by (21c), which follows full exposure to K_θ and then uses V_S . This candidate solves $qV = c - \theta\lambda(a-k) + vV'$ and matches the safe value. Its derivative at K_θ equals $(\theta\lambda a - c)/v$, also the derivative in (32), so it is continuously differentiable at the boundary.

To verify the full action below an interior cutoff, set $F(k) = qV(k) - \theta\lambda k$ and $R = qc + \theta\lambda(v - qa)$. If $r \geq q$, $R > 0$ immediately. If $r < q$, existence of the interior cutoff requires $\theta < cq/\{(q-r)\lambda a\}$, also implying $R > 0$. The full-action equation yields $vF' = qF - R$, with $F(K_\theta) = 0$. Therefore

$$F(k) = \frac{R}{q}\{1 - e^{q(k-K_\theta)/v}\} > 0 \quad \text{for } k < K_\theta. \quad (33)$$

In the full-action region, $C(k) = F(k)$, so full exposure indeed maximizes the HJB expression. Moreover, $vV' = F(k) + \theta\lambda a - c > 0$, verifying the derivative condition used to reduce the action set.

The candidate is bounded, continuously differentiable across the interior boundary, and satisfies (31) in each region and at the boundary. Applying the HJB inequality to $e^{-qt}V(k_t)$ and integrating bounds the objective of every measurable control by the candidate value. The terminal term vanishes as time tends to infinity because $q > 0$ and the candidate is bounded. Policy (21b) attains equality, including the full-exposure continuation at a . This verifies global optimality within problem (19), not merely optimality over the numerical action menu.

Derivation of the policy response

The objective in (28) is strictly convex in b because $\chi > 0$. If $a_0X \leq k$, both terms are minimized at a_0 . Otherwise, in the positive-shortfall region the first-order condition is $\chi(b - a_0) + \nu X(bX - k) = 0$, giving (29). For $X > 0$ and $k < a_0X$, this solution satisfies $k/X < b^* < a_0$, verifying the region and the bounds.

Differentiation with respect to X gives

$$\frac{\partial b^*}{\partial X} = \frac{\nu\{\chi k - 2\chi a_0X - \nu kX^2\}}{(\chi + \nu X^2)^2} < 0,$$

because $k < a_0X$ makes the first two terms strictly negative in sum and the remaining term is nonpositive. The derivative with respect to k is positive as stated. At $X = 0$, there is no incremental exposure and $b^* = a_0$.

Time varying parameters and a conservative bound

Suppose pre-event savings and baseline surplus satisfy $c_t \geq \underline{c} > 0$ and $g_t \geq \underline{g}$ pathwise, with constant ρ . Consider the same stress boundary $x \leq k/a$. Comparison of the maximal-path equations implies that

the constant-parameter path using $(\underline{c}, \underline{g})$ is a lower bound on headroom from the same starting point while the bounds hold. If $-\rho\underline{c} < \underline{g} \leq 0$, equation (10) with those lower bounds is therefore sufficient to complete the transition by H safely on every such path.

This is a sufficient bound and need not be necessary if the worst parameter combination cannot persist. Using average rather than lower-bound savings removes the pathwise guarantee. If $\underline{g} + \rho\underline{c} \leq 0$, the conservative guarantee cannot deliver a self-financing transition from below the target. Replacing a real economic lower bound with a low sample percentile would not establish the guarantee either.

Appendix B computational methods

The Bellman equation and the action menu

For the loss-tolerant problem with $g = 0$, a constant action x chosen at capital k and held over interval Δ gives pre-event capital $k + vxt$. Write $D = (ax - k)_+$, $w = vx$, and $L = \min\{\Delta, D/w\}$ when $w > 0$, with zero penalty when $w = 0$. Exact expected discounted shortfall within the interval is

$$I_\Delta(k, x) = \lambda \left[\frac{D(1 - e^{-qL})}{q} - \frac{w\{1 - e^{-qL}(1 + qL)\}}{q^2} \right]. \quad (34)$$

The corresponding Bellman branch is

$$\frac{cx(1 - e^{-q\Delta})}{q} - \theta I_\Delta(k, x) + e^{-q\Delta} V(\min\{a, k + vx\Delta\}). \quad (35)$$

The state is capped at a because subsequent shortfall is zero and the full-exposure value is $V(a) = c/q$. The within-interval event loss is integrated before this cap; the two steps are consistent because there is no positive shortfall after reaching a .

An additional branch follows the exact continuously adjusted safe policy. Before it reaches a , capital is ke^{rt} , exposure ke^{rt}/a , and interval loss is zero. The script integrates its benefits analytically, splits at the target hit if needed, and uses its exact endpoint for interpolation. This branch keeps a restrictive fixed-action interval from artificially excluding the continuous safe benchmark.

With increasing capital, the continuation state never lies below the starting state. On a capital grid, continuation is a linear interpolation of already solved higher states and possibly the current state. If a branch has reward R , known continuation component C , current-state interpolation weight ω , and discount-survival factor $\beta = e^{-q\Delta}$, its fixed-point value is $(R + \beta C)/(1 - \beta\omega)$. Maximizing these branch values solves the discrete Bellman equation: every denominator is positive, and the largest individual fixed point is exactly the fixed point of the maximum of the affine branches.

The script includes a state-specific safe constant action in addition to the adaptive safe branch. The descending algorithm therefore avoids value-iteration tolerance error. It still has state interpolation, action discretization, and temporal menu error. Joint refinement checks those errors together, rather than proving a separate asymptotic convergence theorem.

The Markov equations

For a short interval dt in state s , capital advances by $v_s dt$, the currency event arrives with probability $\lambda_s dt$, the cost state changes with probability $\mu_s dt$, and discounting contributes $1 - \delta dt$. Keeping first-order terms in the dynamic programming identity for $F_s(k)$ yields

$$0 = v_s F'_s(k) - (\delta + \lambda_s + \mu_s) F_s(k) + \mu_s F_{-s}(k) + \lambda_s (a - k).$$

The terminal conditions are $F_L(a) = F_H(a) = 0$. The code solves this two-equation system backward in k from a to k_0 using an adaptive Runge-Kutta solver with relative tolerance 10^{-10} and absolute

tolerance 10^{-12} . The stationary-weighted value at k_0 gives the exact-model entries in Table 3 up to solver error.

The Monte Carlo check approximates switching and shock arrival on a weekly grid. It uses exact one-step exponential arrival probabilities but applies period-start losses, so it retains a time-step bias. The acceptance check allows five Monte Carlo standard errors plus 10^{-4} in normalized loss for this discretization. That allowance is a numerical check tolerance, not a statistical confidence interval. The ordinary differential equation, rather than the simulation, is the primary estimator of the model expectation.

Replication inventory

Run `python replicate.py` from the replication directory with Python, NumPy, SciPy, and Matplotlib installed. The run writes five figures in PDF and PNG format and eight CSV tables: transition times, recognition rules, heterogeneity, optimal policies, grid convergence, joint risk, parameter-region checks, and the policy phase map. The summary JSON records parameter values, package versions, and numerical checks. No external dataset is required for those original theoretical experiments. The integrated replication also includes the cached CBC, TWSE and Treasury data and the additional scripts described in Appendix C.

The original theoretical outputs have no empirical sampling interpretation except for Monte Carlo sampling uncertainty. The later public-data regressions have their own sampling diagnostics and identification limits. Neither simulated model states nor numerical checks count as additional empirical observations.

Appendix C expanded methods and experiment inventory

Capital transfers and heterogeneous targets

The original dispersion experiment assumes no transfers. A tractable extension makes that restriction explicit. For equally sized institutions with transition targets z_i , define total deficits $D = \sum_i (z_i - k_i)_+$ and surplus $S = \sum_i (k_i - z_i)_+$. Suppose surplus above a donor’s own target can be transferred to any recipient, with fraction $\eta \in [0, 1]$ arriving as usable loss absorption, and injections are allocated after this reallocation. Then the minimum external injection per institution is

$$J(\eta) = \frac{[D - \eta S]_+}{N}. \quad (36)$$

Every recipient deficit must be met by transferred usable surplus or external resources, proving necessity. Transferring up to $\min(S, D/\eta)$ from donor surpluses and allocating the remaining injection attains the bound; for $\eta = 0$, transfer contributes nothing. The extension presumes enforceable transfers before the deadline and no counterparty or timing friction beyond the haircut. It does not prove that actual insurance-group capital is fungible.

The grid uses the same 200 deterministic quantiles as the original distribution experiment, 61 dispersion values, five deadlines and five transfer efficiencies. The 1,525 frontier points are model calculations. A separate comparison holds the marginal distributions of capital and saving rates fixed but reverses their pairing. When weak-capital institutions also have low savings, their targets are more demanding; pairing weak capital with high savings changes the gap even with identical marginal means. The convex-order statement for common targets cannot be generalized to an unrestricted joint distribution.

Deadline	Capital-saving pairing	Transfer efficiency	External budget
3	weak capital / low savings	0%	2.953%
3	weak capital / low savings	100%	1.703%
3	weak capital / high savings	0%	2.106%
3	weak capital / high savings	100%	1.703%
5	weak capital / low savings	0%	1.998%
5	weak capital / low savings	100%	0.088%
5	weak capital / high savings	0%	0.612%
5	weak capital / high savings	100%	0.088%
10	weak capital / low savings	0%	0.946%
10	weak capital / low savings	100%	0.000%
10	weak capital / high savings	0%	0.000%
10	weak capital / high savings	100%	0.000%

Table C1. Hypothetical, equally sized institutions; common marginal capital and saving distributions with different dependence. Pairing affects the budget when transfers are restricted. At full transfer efficiency, the fixed marginal distributions imply the same aggregate budget in both pairings.

Numerical and empirical uncertainty

The original 158 numerical checks are reproduced and retained. They assess internal consistency, including ODE, quadrature, HJB, Bellman and Markov comparisons. They do not establish institutional calibration, originality or causal validity. The review’s checks are also retained; many are source-status

or arithmetic checks rather than independent economic tests. The new continuation calculation uses 801 capital nodes and 101 constant actions, with an exact adaptive safe branch at a weekly step. Saved sensitivity surfaces contain 32,160 horizon-capital records and 36,381 reward-capital records. Those are model states, not additional financial observations.

The daily empirical covariance is Newey-West with ten observed-row lags and a finite-sample factor. Weekly regressions use four lags, monthly accounting regressions six, and the mean forecast-error difference twenty. Normal-approximation intervals can be unreliable under strong nonstationarity or exceptional jumps. The block comparisons, year omissions and regime-specific estimates are reported to expose that dependence. No specification is selected as a causal model because its interval excludes zero.

Monthly records include 18 missing foreign-asset entries, 73 missing book-equity entries, seven missing other-liability entries and eight missing observed contract-liability entries. Missing values are not converted to zeros. The foreign-asset translation decomposition is only available where adjacent stocks and CBC exchange-rate endpoints are available. The US rate source is Treasury par yields, which differ from duration-matched bond returns and from the actual bonds on insurer balance sheets.

The primary protocol is retrospective. Both the May episode and the earlier five-parent findings were known when the extension was chosen. Historical source vintages downloaded now are not a reconstruction of every revision available to an investor on each original date. Forecast information is chronologically lagged, but the evaluation is not claimed to have been prospectively preregistered. These restrictions apply even when the row count is large.

Selected parent universe and coverage

Ticker	Parent	First observation	Last observation	Usable return days
2880	Hua Nan	2010-01-04	2026-10-22	4099
2881	Fubon	2010-01-04	2026-10-22	4095
2882	Cathay	2010-01-04	2026-10-22	4093
2883	KGI / China Development	2010-01-04	2026-10-22	4100
2884	E.SUN	2010-01-04	2026-10-22	4094
2885	Yuanta	2010-01-04	2026-10-22	4094
2886	Mega	2010-01-04	2026-10-22	4097
2887	Taishin / TS	2010-01-04	2026-10-22	4086
2888	Shin Kong	2010-01-04	2025-07-11	3790
2889	Waterland	2010-01-04	2026-10-22	4093
2890	SinoPac	2010-01-04	2026-10-22	4097
2891	CTBC	2010-01-04	2026-10-22	4093
2892	First	2010-01-04	2026-10-22	4100
5880	Taiwan Cooperative	2011-12-01	2026-10-22	3624

Table C2. Source observations and usable clean-return dates. A larger count in this table is not the number of independent currencies or institutions supplying economic-capital data. Delisted and prelisting periods remain missing.

Original comparison and additional exposure checks

Original basket specification	Dates	Coefficient	95% HAC interval
2010 onward	4024	-0.316	[-0.507, -0.125]
2010–2019	2411	-0.373	[-0.613, -0.134]
2020–2024	1187	-0.054	[-0.405, 0.296]
2025 onward	426	-0.349	[-0.864, 0.167]
excluding May stress	4015	-0.274	[-0.454, -0.093]
reference returns	4116	-0.313	[-0.492, -0.134]
prior-US yield controls	4022	-0.298	[-0.501, -0.095]

Table C3. The original three-parent comparison, expanded in time. Its 2020-onward historical result remains separately reproducible from the earlier analysis package; the full-period specification here deliberately changes the time window.

The longer price history also supplies the previous close needed to construct the first return in January 2020. An expanded-history estimate restricted to 2020 can consequently have one more usable observation than the original file that began in January 2020. This is a documented history-boundary change, not an imputed price.

Sample	Positive move	Coefficient	95% HAC interval
all	app	-0.284	[-0.591, 0.024]
all	dep	0.262	[-0.040, 0.565]
excluding May	app	-0.148	[-0.385, 0.089]
excluding May	dep	0.296	[-0.006, 0.599]

Table C4. Separate positive appreciation and depreciation magnitudes. A positive depreciation regressor has the opposite currency sign from a positive appreciation regressor. These are descriptive slope comparisons, not isolated directional shocks.

Complete experiment crosswalk

ID	Experiment	Retained contributing work
E1	Safe transition and viability	Closed forms, ODE checks, deficit and capital grid
E2	Recognition and first-event shortfall	Four policies, quadrature, 500,000 event draws
E3	Deadline funding and capital distribution	Original dispersion grid; transfer and joint-target extension
E4	Loss-tolerant optimum	Exact proof, independent Bellman, refinement, parameter regions
E5	Savings and event dependence	Markov backward equations and 150,000 common paths
E6	Reward definition	Original four rewards and dense reward-capital grid
E7	Horizon and continuation	Original five horizons; 80-horizon continuation comparison
E8	Public insurance balance sheet	472 monthly source rows, identity and definition checks
E9	Equity exposure	Original and expanded panels; rates, asymmetry, block and era checks
E10	Event and timing evidence	Original events, placebo ranks, unavailable-window audit
E11	Chronological prediction	Original conditional diagnostic; annual expanding evaluation
E12	Marked-account historical replay	166 overlapping five-year paths; 2025 path
E13	Rebalancing and costs	190 overlapping three-year histories; 27 full-factorial cases
E14	Conditional forward environments	Two historical pools; three block lengths; 6,000 paths per pair
E15	Named stresses and exit liquidity	Five specified paths; proportional closeout correction

Table C5. Complete experiment inventory. The first five families are the original theoretical experiments; the later families retain the critical review and the expanded public-data evidence.

All original five experiment families and the later review experiments are included. Additional panels, sensitivity tables and uncertainty checks are grouped by their substantive question. The full source, raw-data cache, clean tables, checks and vector figures are provided together. Simulation samples, deterministic grids, rolling windows and independent observation dates retain separate meanings throughout.

The ten simulated expert perspectives from the critical review remain in the replication as a reasoning and revision record. They are not interviews, human peer-review reports or independent endorsements. The integrated paper incorporates their substantive concerns through the reward, continuation, liquidity, measurement and identification sections rather than reproducing ten voices in its argument.

Transition-gap revision: additional experiment inventory

Family	Completed calculation	Interpretation
E16	Joint asset, liability and forward reconciliation	Exact illustrative accounting; no estimated Taiwan surplus
E17	561 funding-frontier cases, 108 sensitivity cases and matched dated paths	Synthetic conditional feasibility; ordinary full liquidity ladders can contain the same information
E18	Six insurer-report periods, three bridges each	Audited arithmetic of reported totals; economic gaps remain unidentified
E19	Eight main policy cases, two net-flow cases and nine opening-resource cases, independent common-path checks	Optimum in a discrete-action model with stipulated flows and penalties; no default forecast
E20	Derivative-maturity, nominal and gross/net accounting audits	Reported stocks and buckets do not identify gross payment requirements or hedge ratios
E21	Asset sale, secured-credit and coarse-calendar funding cases	The synthetic timing gap can disappear when existing assets can be monetized

Table G9. New work alongside the fifteen earlier families in Table C5. The predictive protocol and proposed causal design are not counted as completed empirical experiments. Exploratory interpolation results are retained as numerical diagnostics and superseded by the exact lattice for principal conclusions.

References

- Australian Accounting Standards Board. 2020. *AASB 7: Financial Instruments—Disclosures*, September 2020 compilation, Appendix B11D–B11F. [Public primary text](#). Used for the net-versus-gross cash-flow distinction, not as a statement of Taiwanese law.
- Bai, Jennie, Arvind Krishnamurthy, and Charles-Henri Weymuller. 2018. “Measuring Liquidity Mismatch in the Banking Sector.” *Journal of Finance* 73(1): 51–93. [Author institution record](#); [NBER working paper 22729](#).
- Calvo, Guillermo A., and Carmen M. Reinhart. 2002. Fear of Floating. *Quarterly Journal of Economics* 117 (2): 379–408. [doi:10.1162/003355302753650274](#).
- Cathay Life Insurance Co., Ltd. and Subsidiaries. 2024 annual and 2025 first-half consolidated financial statements. [Issuer report index](#). Exact report URLs, PDF pages and SHA-256 hashes: `data/institution_source_manifest.json` and `data/institution_fields.csv`.
- Central Bank of the Republic of China (Taiwan). 2026a. *Financial Stability Report*, May, Issue 20, especially pp. 65–68 and Appendix Table 2. [Official report](#).
- Central Bank of the Republic of China (Taiwan). 2026b. Statement concerning questions about life insurers’ foreign exchange accounting reform. January 30. Chinese language primary statement. [Official statement](#).
- Central Bank of the Republic of China (Taiwan). 2026c. Assets and Liabilities of Life Insurance Companies, monthly open data, May 1987–August 2026. [Official CSV](#). Definitions and changes checked against [Financial Statistics Monthly Appendix 8](#).

Central Bank of the Republic of China (Taiwan). 2026d. Daily interbank NTD/USD closing observations. [Official open-data API](#). The archived extract ends October 22, 2026.

Du, Wenxin, and Amy W. Huber. 2026. Dollar Asset Holdings and Hedging around the Globe. *Review of Financial Studies*, advance article hhag067, published July 27. [doi:10.1093/rfs/hhag067](#).

Ellul, Andrew, Chotibhak Jotikasthira, Christian T. Lundblad, and Yihui Wang. 2015. Is Historical Cost Accounting a Panacea? Market Stress, Incentive Distortions, and Gains Trading. *Journal of Finance* 70 (6): 2489–2538. [doi:10.1111/jofi.12357](#).

European Insurance and Occupational Pensions Authority. 2021. “Methodological Principles of Insurance Stress Testing: Liquidity Component.” January 26. [Official publication record](#).

European Insurance and Occupational Pensions Authority. 2024. “Insurance Stress Test 2024.” [Official exercise and results](#).

Executive Yuan Gazette. 2025. Foreign-exchange valuation reserve order effective June 30, 2025, published July 1. Chinese primary legal record. [Official Gazette](#).

Farhi, Emmanuel, and Jean Tirole. 2012. Collective Moral Hazard, Maturity Mismatch, and Systemic Bailouts. *American Economic Review* 102 (1): 60–93. [doi:10.1257/aer.102.1.60](#).

Financial Supervisory Commission. 2025. Proposed amendments to the Regulations Governing the Preparation of Financial Reports by Insurance Enterprises concerning foreign exchange gains and losses. December 23; English page updated January 2, 2026. [Official release](#).

Financial Supervisory Commission. 2026a. Regulations Governing the Preparation of Financial Reports by Insurance Enterprises. Article 12 currency translation proviso. Consolidated text accessed October 26, including prospective 2028 amendments. Chinese language legal text. [Official regulation](#).

Financial Supervisory Commission. 2026b. Directions concerning life insurers’ foreign exchange valuation reserves. Amendment of February 13, especially paragraphs 3 and 7–10. English description of Chinese language legal text. [Official rule](#).

Froot, Kenneth A., David S. Scharfstein, and Jeremy C. Stein. 1993. “Risk Management: Coordinating Corporate Investment and Financing Policies.” *Journal of Finance* 48(5): 1629–1658. [NBER working paper 4084, 1992](#).

Fubon Financial Holdings. 2025. May 2025 earnings results, June 10. [Primary issuer release](#). Contextual issuer reporting; not an independently reconstructed economic-capital account.

Fubon Life Insurance Co., Ltd. and Subsidiaries. 2024 annual and 2025 first-half consolidated financial statements. [2024 official index](#); [2025 official index](#). Exact report provenance is in the institutional source manifest.

Insurance Bureau. 2025. “Financial Overview of Insurance Industry as of May 2025.” June 26. [Official release](#).

KGI Life Insurance Co., Ltd. 2024 annual and 2025 first-half financial statements. [Official report index](#). Exact report provenance is in the institutional source manifest.

McGuire, Patrick, Ilhyock Shim, Hyun Song Shin, and Vladyslav Sushko. 2021. Outward Portfolio Investment and Dollar Funding in Emerging Asia. *BIS Quarterly Review*, December. [Official article](#).

- Rampini, Adriano A., Amir Sufi, and S. Viswanathan. 2014. Dynamic Risk Management. *Journal of Financial Economics* 111 (2): 271–296. [Published paper at author’s institution](#).
- Sen, Ishita. 2023. Regulatory Limits to Risk Management. *Review of Financial Studies* 36 (6): 2175–2223. [doi:10.1093/rfs/hhac083](https://doi.org/10.1093/rfs/hhac083).
- Shin, Hyun Song, Philip Wooldridge, and Dora Xia. 2025. US Dollar’s Slide in April 2025: The Role of FX Hedging. *BIS Bulletin* 105, June 20. [Primary research](#).
- Taiwan Stock Exchange. 2026. Daily individual-security and TAIEX observations, queried by month. [STOCK DAY endpoint example](#); [FMTQIK index endpoint example](#). The replication manifest records all requested months, tickers and no-data responses. [Operating Rules Article 3](#) specifies the regular equity trading hours.
- U.S. Department of the Treasury. 2026. *Macroeconomic and Foreign Exchange Policies of Major Trading Partners of the United States*. July, Taiwan section and Box 1. [Official report](#).
- U.S. Department of the Treasury. 2026b. Daily Treasury Par Yield Curve Rates, annual CSV files for 2007–2026. [Official data page](#). Quote dates are matched to Taiwan dates explicitly; yields are not bond total returns.
- Wang, Bo-Zhuan. 2026. Foreign Exchange Risk Management and Financial Resilience in Taiwan Life Insurance Industry: The Effects of FX Accounting Reform and FX Valuation Reserve Mechanisms. Thesis repository record, National Chengchi University. Repository abstract consulted; full thesis not independently replicated. [University record](#).