Economica Obscura

Institutions, Archives, and Policy Learning
Online Appendix

Chae-Yeon Xon
Yonsei University

September 2026

Abstract

This online appendix contains dynamic and content-choice extensions, state-persistence calculations, posterior calculations, data construction, and predictive diagnostics for the paper. The core model derivations and proofs accompany the main paper. The sections here document additional specifications, the U.S. bank posterior, forecast construction, allocation rules, reform cohorts, and further predictive diagnostics.

JEL Codes: C11, C13, C18, D83, E61

Keywords: identification, information production, institutional adjustment, policy learning, experimental design

Yonsei University. Email: eunixon@yonsei.ac.kr. This draft is circulated for discussion; comments and suggestions are welcome.

Contents

Appendix A. Sequential Reporting and Record Value

A.1. Proof of Theorem A2

For two-dimensional vectors, the determinant expansion of $J_0 + \tau \sum_s e(\delta_s)e(\delta_s)'$ is the sum of the initial determinant, the cross-products with J_0, and the squared areas spanned by pairs of new vectors. The cross-products equal $\tau \sum_s \{b \cos^2 \delta_s + a \sin^2 \delta_s\}$; each squared area equals $\tau^2 \sin^2(\delta_r - \delta_s)$. This proves (A23). The trace equals S_t, so $\operatorname{tr}(J_t^{-1}) = S_t / \det J_t$.

Let $E_t(\delta) = \det J_t - K_t \geq 0$ and let $E_t^{(2)}(\delta)$ replace every squared sine in (A23) by its squared argument. Then $0 \leq E_t \leq E_t^{(2)}$, and

$$F_H(\delta) - F_H(0) = \frac{1}{4}\delta' A_H \delta - \frac{\chi}{2} \sum_{t=1}^{H} \frac{\eta^{t-1} S_t E_t(\delta)}{K_t[K_t + E_t(\delta)]} \geq \frac{1}{4}\delta'(A_H - 2\chi B_H)\delta. \tag{A1}$$

The matrix A_H is positive definite. If $\chi < \chi_H^*$, the last quadratic form is positive for every nonzero δ. At equality, consider a nonzero vector in its null space. It has $\delta' B_H \delta > 0$ and hence $E_t^{(2)} > 0$ for some t. For that t, $E_t/(K_t + E_t) < E_t^{(2)}/K_t$: when $E_t = 0$ the right side is positive, and when $E_t > 0$ the denominator gives strict inequality. Thus the original objective difference is positive. This also establishes uniqueness at the threshold. If $B_H = 0$, information has the same value along every path in the only possible case, $H = 1$ and $a = b$, and the quadratic cost selects zero.

For $\chi > \chi_H^*$, choose z with $z'(A_H - 2\chi B_H)z < 0$. Expansion at ϵz gives $F_H(\epsilon z) - F_H(0) = \epsilon^2 z'(A_H - 2\chi B_H)z/4 + O(\epsilon^4)$, which is negative for sufficiently small positive ϵ. Positive definite quadratic costs make the continuous objective coercive, so a minimizer exists and has a nonzero angle. For $a = b$ and $H \geq 2$, the pair $(1, 2)$ gives a positive quadratic form for a path with $z_1 \neq z_2$. The stated $H = 1$ formula follows by substituting $A_1 = \rho + \kappa$ and $B_1 = (a + b + \tau)\tau(a - b)/[b(a + \tau)]^2$.

A.2. Proof of Corollary A2

With $J_0 = jI_2$, the first posterior inverse trace is independent of δ_1. At the second date it equals $S/[K + \tau^2 \sin^2(\delta_2 - \delta_1)]$. Put $z = \delta_2 - \delta_1$. Four times the discounted operating and adjustment cost equals

$$h\delta_1^2 + \eta\rho(\delta_1 + z)^2 + \eta\kappa z^2 = (h + \eta\rho)\left(\delta_1 + \frac{\eta\rho z}{h + \eta\rho}\right)^2 + \eta c z^2. \tag{A2}$$

Minimizing the square gives the angles in (A27). After removing the first policy loss and dividing the remaining objective by η, the problem is

$$\min_{z \in \mathbb{R}} \left\{ \frac{c}{4} z^2 + \frac{\chi S}{2(K + \tau^2 \sin^2 z)} \right\}. \tag{A3}$$

An information-equivalent representative with $|z| \le \pi/2$ has the smallest cost. The objective difference from zero is $cz^2/4 - \chi S\tau^2 \sin^2 z/[2K(K + \tau^2 \sin^2 z)]$. The argument of Theorem 2 therefore gives the threshold, the unique positive root above it, the reflection, and the limiting angles. In particular, $\sin(2z)/z$ strictly decreases on $(0, \pi/2)$ and the denominator in (A27) increases there. Finally, $\partial c/\partial\kappa = 1+\eta\rho^2/(h+\eta\rho)^2 > 0$ and $\partial c/\partial\eta = -\rho^2 h/(h+\eta\rho)^2 < 0$. The remaining factors of the threshold are constant in these comparisons.

A.3. Proof of Theorem A3

At ζ_0, Corollary A2 reduces the objective to the completed square in (A2) and the scalar problem (A3). For $\chi > \chi_2^*$, its two solutions $\pm z^0$ lie in the interior of $(-\pi/2, \pi/2)$ and the derivative in the stationary equation crosses zero once on the positive branch. The scalar second derivative is positive at z^0. The completed square contributes positive curvature in the remaining direction, so the Hessian of F_2 is positive definite at δ^0 and $-\delta^0$.

The first-order condition $\nabla_\delta F_2(\delta; \zeta) = 0$ is continuously differentiable in angles and positive parameters. The implicit function theorem gives two continuously differentiable stationary branches through the benchmark minima and a Lipschitz bound of the form in (A28) after shrinking the parameter neighborhood. Positive-definite Hessians persist in that neighborhood.

Quadratic operating and adjustment costs make every common sublevel set lie in a compact angle rectangle for parameters near ζ_0. At the benchmark, continuity and uniqueness in Corollary A2 give a positive objective gap between neighborhoods of the two minima and the remainder of this rectangle. Uniform continuity preserves part of that gap for nearby parameters. Every global minimizer therefore belongs to one of the two stationary branches. Since both coordinates of δ^0 are separated from zero, the Lipschitz bound preserves their opposite signs after one final reduction of the neighborhood.

A.4. Numerical Map for Theorem A3

A numerical map varies a/b and d_1/d_2 over 31 logarithmically spaced values from $\exp(-0.55)$ to $\exp(0.55)$, holding $ab = d_1 d_2 = 1$. At each of 961 points it minimizes the two-angle objective from 25 initial values on $[-1.35, 1.35]^2$. The reports straddle zero at 97.09 percent of the full grid and every point with both ratios in $[0.8, 1.25]$. Within that rectangle the smaller absolute angle is at least 0.115 radians.

A.5. Proof of Lemma 4

The Gaussian conditioning formula gives $P_t^+ = P_t - P_t ss' P_t/(\nu + s' P_t s)$. Future independent innovations contribute the same covariance under the baseline and augmented archives. The covariance difference equals $A^h P_t ss' P_t (A^h)'/(\nu + s' P_t s)$. Taking one half of its trace

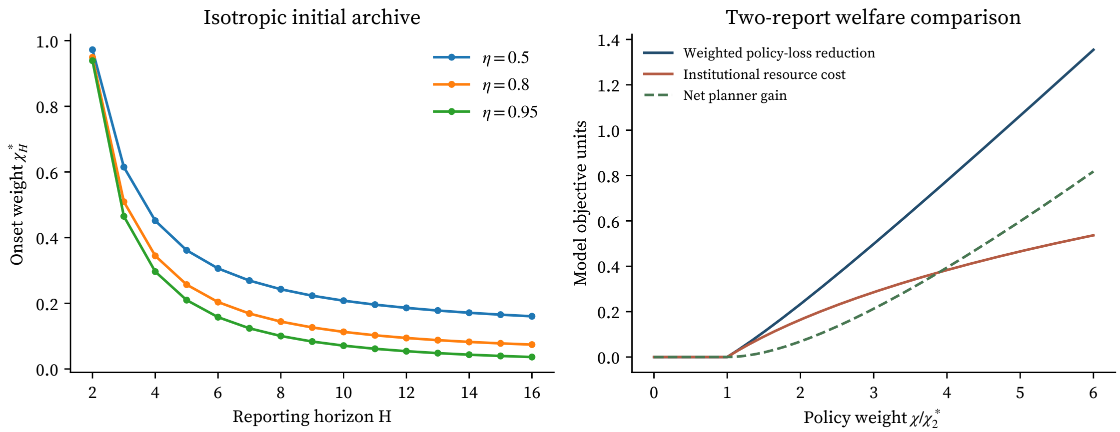

FIGURE A1. Reporting horizon and the welfare wedge

The left panel evaluates the threshold for $a = b = 1$, $\tau = 3$, $\rho = 1$, $\kappa = 2$, $H = 2, \ldots, 16$, and the displayed discount factors. These horizon comparisons are numerical. The right panel uses two dates and $\eta = 0.95$ to decompose equation (24). Costs and benefits are measured in model units.

against D proves (15). Its numerator equals the squared norm of $D^{1/2}A^h P_t s$, which gives the zero-value condition. The covariance may be singular; the argument only uses positivity of the report's predictive variance.

A.6. Proof of Corollary A1

For $B_H \neq 0$, $\chi_H^* = \min_{z:z'B_H z=1} z' A_H z/2$. Define $W_\eta = \mathrm{diag}(1, \eta, \ldots, \eta^{H-1})$ and the difference operator $(\Delta z)_t = z_t - z_{t-1}$ with $z_0 = 0$. The cost matrices satisfy

$$A_H = \rho W_\eta + \kappa \Delta' W_\eta \Delta. \tag{A4}$$

The square matrix Δ is invertible. Increasing either cost adds a positive-definite matrix E to A_H. The normalization $z'B_H z = 1$ implies $z'Ez \geq \lambda_{\min}(E)/\lambda_{\max}(B_H) > 0$, so the minimized quotient strictly increases. Joint scaling of ρ and κ scales A_H and leaves B_H fixed, which proves homogeneity.

Appendix B. A Report-Content Cutoff after a Precision Update

B.1. Proof of Theorem A1

The covariance formulas in (A16) follow from conditioning a Gaussian vector on $Y = \theta_1 + \epsilon$. Conditional on $Y = y$, T is Gaussian with mean $\alpha + \lambda k y$ and variance $\sigma_T^2 + \lambda^2 p_1$. This gives $\pi(y)$ and shows that it is strictly increasing in y.

Applying the rank-one covariance update to the second report gives (A18). The variance reduction from an e_i report is $A(p_i)e_ie_i'$, where $(p_1, p_2) = (p_1, s_2)$. Hence

(A5)
$$R_1(y) - R_2(y) = \frac{1}{2}\left[d_2\{\pi(y)\}A_2 - d_1\{\pi(y)\}A_1\right] = \frac{1}{2}\left[d_H A_2 - d_L A_1 - \Delta_d(A_1 + A_2)\pi(y)\right].$$

The move is selected when $c < \chi\{R_1(y) - R_2(y)\}$. Solving for $\pi(y)$ gives (A19). Strict monotonicity of $\pi(y)$ and inversion of the standard normal distribution give (A20).

For every realization, the adaptive plan takes the smaller of the two conditional costs. Taking expectations gives a value weakly below the expectation under either fixed action and therefore gives (A21). If each action is strictly preferred on an event with positive probability, each fixed action incurs a strictly larger expected cost on one of those events.

Appendix C. Proof for a Changing State

C.1. Proof of Proposition 2

Substituting $A^h = \phi^h I_2$ into (15) gives $\Delta_h = \phi^{2h}\Delta_0$. Summing the discounted geometric sequence gives the second expression in (16).

For the specialized stream, write $P_{T+h} = \mathrm{diag}(p_{1,T+h}, p_{2,T+h})$. Diagonality is preserved because the state innovation covariance and the inherited covariance are diagonal and the report measures the first coordinate. Prediction followed by measurement gives

(A6)
$$s_{i,h} = \phi^2 p_{i,T+h-1} + q, \qquad p_{1,T+h} = \frac{s_{1,h}}{1 + \tau s_{1,h}}, \qquad p_{2,T+h} = s_{2,h}.$$

The second-coordinate recursion is affine with slope ϕ^2, so it converges to $q/(1-\phi^2)$ and differences between two inherited values equal ϕ^{2h} times their initial difference. For the first coordinate, define $f(p) = (\phi^2 p + q)/\{1 + \tau(\phi^2 p + q)\}$. On $[0, \infty)$, $0 \le f'(p) = \phi^2/\{1 + \tau(\phi^2 p + q)\}^2 \le \phi^2 < 1$. The contraction mapping theorem gives a unique fixed point and the first bound in (18). Solving $p = f(p)$ gives the positive root in (17); direct substitution gives the displayed value at $\phi = 0$. Posterior decision loss is $\mathrm{tr}(DP)/2$. Since both covariance paths are diagonal and $D_{11}, D_{22} \ge 0$, applying (18) gives the final bound.

Appendix D. Posterior Construction from Bank Archives

D.1. Likelihood and Posterior

Consider retained reports $y_1, \ldots, y_m$ at dates $s_1 < \cdots < s_m$, with gaps $d_k = s_k - s_{k-1}$ and increments $z_k = y_k - y_{k-1}$. Conditional on the first report and a bank-specific pair (b, q), the likelihood is proportional to $q^{-(m-1)/2}\exp\{-\sum_{k=2}^{m}(z_k - d_k b)^2/(2qd_k)\}$. The level prior is diffuse, $b \mid q \sim N(0, q/\lambda_0)$, and $q \sim \mathrm{IG}(a_0, \beta_0)$ under the density convention $q^{-a_0-1}\exp(-\beta_0/q)$. Put $\beta_0 = (a_0-1)q_0$, so the prior mean innovation variance is q_0.

Let $S = s_m - s_1$, $Z = y_m - y_1$, and $U = \sum_{k=2}^{m} z_k^2/d_k$. Completing the quadratic gives

$$\text{(A7)} \qquad \lambda = \lambda_0 + S, \qquad \widehat{b} = \frac{Z}{\lambda}, \qquad a = a_0 + \frac{m-1}{2}, \qquad \beta = \beta_0 + \frac{1}{2}\left(U - \frac{Z^2}{\lambda}\right).$$

The posterior is $b \mid q, H \sim N(\widehat{b}, q/\lambda)$ and $q \mid H \sim \mathrm{IG}(a, \beta)$. Thus $\widehat{q} = \beta/(a-1)$ and $\mathrm{Var}(b \mid H) = \widehat{q}/\lambda$. For a target r quarters after the latest report, its predictive distribution has $2a$ degrees of freedom, location $y_m + r\widehat{b}$, and squared scale $(\beta/a)(r + r^2/\lambda)$. Its variance is (A29) with $r = d_{it} + h$. The 90 percent intervals use the corresponding Student-t quantile.

At a current report, the level is known and its posterior variance is zero. Drift and future innovations retain positive variance. At age $d > 0$, the covariance of the current level-and-drift state is

$$\text{(A8)} \qquad P_t = \widehat{q}\begin{pmatrix} d + d^2/\lambda & d/\lambda \\ d/\lambda & 1/\lambda \end{pmatrix}.$$

The transition matrix is $A = \left(\begin{smallmatrix} 1 & 1 \\ 0 & 1 \end{smallmatrix}\right)$, with conditional innovation covariance $q\,\mathrm{diag}(1,0)$. Equation (A8) describes a scale mixture of Gaussian posteriors. The full archive state also contains the posterior mean and inverse-gamma parameters. Conditional drift precision j^b remains defined when the current level has been observed exactly.

D.2. Expected Gain from an Intermediate Report

Before a current report is acquired, its predictive variance is $\widehat{q}(d + d^2/\lambda)$ and its covariance with the target h quarters later is $\widehat{q}\{d + d(d+h)/\lambda\}$. Its posterior-mean update is linear, with a coefficient independent of q. The law of total variance therefore gives expected MSE reduction

$$\text{(A9)} \qquad \frac{\widehat{q}\{d + d(d+h)/\lambda\}^2}{d + d^2/\lambda} = \widehat{q}\left\{d + h + \frac{(d+h)^2}{\lambda} - h - \frac{h^2}{\lambda + d}\right\}.$$

Integrating the conditional Gaussian calculation in Lemma 4 gives the same formula. Expected posterior q after the new report equals its current posterior expectation, and the updated drift precision factor is $\lambda + d$. At $d = 0$, the report repeats a known exact level and the gain is zero. The March and September semiannual archives have positive age.

D.3. Annual Prior Selection and Predictive Diagnostics

Prior selection uses $\lambda_0 \in \{4, 16, 64, 256, 1024\}$, $a_0 \in \{3, 6, 12\}$, and $q_0 \in \{0.25, 1, 4, 16, 64\}$. The 75 tuples are compared by mean negative predictive log density in year $y - 1$, averaged equally across available-twelve and semiannual-twelve archives. Ties use grid order. The selected tuple is held fixed for target year y, and bank posteriors use their retained origin records. Validation outcomes are dated at or before the first evaluation origin. Selected a_0 is 3 in every year. Selected q_0 is 1 in 2018 and 0.25 in the remaining years, at the lower grid boundary.

The model uses unbounded Gaussian increments for a bounded share. Predictive means are untrimmed in the posterior-state results so that loss and covariance calculations use the same specification. Coverage is evaluated from Student-t predictive intervals. Older histories can raise $\widehat{q}$ when they contain larger past changes. Pointwise posterior variance can consequently rise when those records enter, as in Table A7. Future-error diagnostics assess whether estimated dispersion corresponds to subsequent variation.

The same-current-report comparison uses main-sample rows with an observed origin ARM value and twelve available observations. Four-, eight-, and twelve-record histories retain the same current record. For variance calibration, origin-specific current-composition deciles and asset quintiles define joint cells. After subtracting cell means from squared error and predicted variance, ordinary least squares estimates their slope. Gain calibration uses latest semiannual composition deciles and asset quintiles within origin. Cell counts enter the residual degrees-of-freedom correction. Bank and target-date score totals give clustered slope variances. Negative estimated two-way variance is flagged and its interval left unreported.

Annual tuning and evaluation use chronological splits: prior parameters are selected from the preceding year, target-year posteriors use records available through each origin, and downstream forecasts use earlier target years. The diagnostic calculations use the same dated information sets.

Appendix E. Archive Score and the Report Budget

Under equal item costs and additive squared-error targets, the criterion maximizes $\sum_i a_i g_i$ over $a_i \in \{0, 1\}$ subject to $\sum_i a_i \le k_t$. Its solution sorts g_i, with certificate number breaking ties. Uniform expected allocation assigns probability k_t/n_t to each eligible report within an origin. Heterogeneous monetary costs replace the constraint by $\sum_i c_i a_i \le B_t$ and require the corresponding knapsack solution. Each evaluated allocation starts from its inherited semiannual history and holds the bank, realized report, and later outcome fixed across rules.

An accounting decomposition writes (34) as $\widehat{q}d + \widehat{q}\{(d+h)^2/\lambda - h^2/(\lambda+d)\}$. The first term is the value of observing accumulated innovations when drift is known; the second comes from uncertainty about drift. Both use the variance learned from the archive. Over the observed-origin sample, the second term accounts for 6.26 percent of total predicted gain. This decomposition explains why innovation scale provides a demanding benchmark for the mixed-state allocation rule.

Appendix F. Empirical Data and Sample Construction

The downloaded financial panel covers all 56 quarters from 2012 through 2025. Its 89 requested fields include identifiers and reporting metadata. Form 041 and 051 domestic commercial banks account for 268,189 source rows. The API field `LNRERSF1` records adjustable-rate first-lien mortgages; `LNRERSFM` records the corresponding first-lien total. Monetary fields are in thousands of dollars. The six mortgage maturity or repricing buckets plus first-lien nonaccrual mortgages agree with the total within four thousand dollars throughout this source population. The comparison allows for rounding of separately reported entries.

The other-loan maturity reconciliation carries a flag for 11,436 source rows after including the gross-loan and unearned-income bridge. The strict sample excludes flagged origin or target observations. Aggregate interest-rate derivatives, `RT`, are included in common controls. Component derivatives fields have reporting gaps, so their values remain outside the input set. Scheduled zeros in the ARM item are retained when their total is positive. All March and September form-051 ARM entries are masked, including the 259 positive entries found in the domestic commercial source population.

Bank histories use dated FDIC certificates and RSSD identifiers. A merger, an identifier change, or a quarterly gap starts a new archive spell. Main-sample origins require three available ARM records, a June or December record, the common controls, and both future outcomes. The archive spell must continue through the income and ARM target dates. Quarterly flow reconciliation flags at the origin or either target and unresolved institution metadata are excluded. These target-date conditions describe the evaluated sample after outcomes are observed. The broad sample retains event and accounting flags subject to observed labels and origin controls; the strict sample applies the additional other-loan accounting filter.

The FRED inputs are the available daily observations of `DFF`, `DGS2`, and `DGS10` within each quarter. Changes compare the current quarterly mean with its one- and four-quarter lags. The first four quarters use zero for an unavailable four-quarter rate change. All subsequent changes are dated before or at the forecast origin. The FDIC data use the August 28, 2026 download vintage.

TABLE A1. Evaluation samples and target dates

Period	Sample	Rows	Banks	ARM dates	Income dates
2016	Main	19,480	5,070	2	4
	Other-loan check	18,374	4,995	2	4
	Broad	19,820	5,129	2	4
	Observed odd quarter	9,799	5,070	2	2
2022–2025	Main	59,892	4,096	8	16
	Other-loan check	55,349	4,087	8	16
	Broad	60,542	4,098	8	16
	Observed odd quarter	7,085	1,018	8	8
2016–2025	Main	165,049	5,294	20	40
	Other-loan check	148,956	5,282	20	40
	Broad	167,473	5,301	20	40
	Observed odd quarter	29,588	5,166	20	20

Rows are bank-origin observations with both targets. The main sample conditions on stable spells and the stated accounting checks. The broad sample retains event and accounting flags; the other-loan sample adds its reconciliation filter. Observed odd quarters are March and September origins with an actual scheduled ARM value.

Appendix G. Forecasting and Selection Algorithms

Let x_{it}^{s} denote the common and archive features for scenario s. The scenarios retain up to twelve available reports, twelve semiannual reports, one available report, four available reports, eight available reports, four reports after discarding one recent report, four after discarding two, and the semiannual archive with the current report restored. Training observations with shorter histories retain their available records and record the resulting count. The dedicated record-length and recency evaluations require the specified counts in every compared scenario.

For each outcome, the ridge coefficient minimizes

$$\frac{1}{8N_{\text{train}}}\sum_{s=1}^{8}\sum_{it\in\text{train}}(Y_{it}-a-x_{it}^{s\prime}b)^2+\lambda b'b. \tag{A10}$$

The intercept is unpenalized. Financial ratios and current and lagged income features are clipped at their training 0.5 and 99.5 percentiles. Missing lagged entries use training medians with a lag-availability indicator. Features are standardized using training means and standard deviations. Rate levels and changes retain their observed ranges, and responses retain their observed values. The penalty grid is $\{10^{-4},10^{-3},10^{-2},10^{-1},1\}$. Validation selects the smallest average MSE over the available-report and semiannual scenarios, with equal weights. Ties select the larger penalty. The selected penalty is used after re-estimation and preprocessing on all preceding target years. ARM predictions are restricted to $[0, 100]$.

Common controls contain log assets; mortgage, capital, deposit, securities, loan, nonaccrual, and aggregate interest-rate derivative balances scaled by assets; six mortgage maturity shares; six other-loan maturity balances; eight deposit maturity balances; twelve securities maturity balances; current and lagged income and its change; current rate levels and dated changes; the Treasury slope; quarter indicators; and the known ARM horizon. ARM archives contribute the latest percentage share, mean, standard deviation, age, count, span, and slope. The latest and average share interact with the current mortgage-to-asset ratio, and the latest exposure also interacts with the federal funds rate and its quarterly change.

The persistence forecast uses the last observed ARM share and current quarterly net interest income divided by current assets. The Gaussian comparator first regresses observed current ARM shares on common controls, using a fixed ridge penalty of 0.01. It filters residuals with a scalar random-walk state and temporary observation variation. The state-innovation to observation-variance ratio is selected from $\{0.001, 0.01, 0.1, 1, 10\}$ using validation forecast MSE, averaged over the two schedules. Prior residual variance is ten observation-variance units. Each retained archive is processed in date order, and gaps enter the state-variance update. The ARM forecast combines the filtered residual with the current common-control prediction and is clipped to $[0, 100]$. A separate income ridge regression uses the common controls and this filtered exposure. Its penalty uses the same validation grid. This is a statistical forecasting approximation; the residual observation variance represents temporary composition variation.

The auxiliary mortgage-total forecast uses common controls and a validation-selected ridge penalty. Its target is next-quarter first-lien mortgages divided by origin assets in basis points. Its input set is common to the reporting scenarios, so the schedule exercise supplies the same total forecast in each scenario.

For allocation, define validation gain $g_{it} = (Y_{it} - \widehat{f}_{it}^{S})^2 - (Y_{it} - \widehat{f}_{it}^{S+})^2$. A ridge regression with penalty 0.1 projects g_{it} on the semiannual feature vector in the preceding validation year, using that year's preprocessing. The score is fitted separately for each target. Selection ranks the resulting score within the evaluation origin, breaking ties by certificate. Asset selection uses the same tie breaker. A fixed seed generates uniform draws. For k reports and n eligible banks, the expected uniform loss is $n^{-1} \sum_i [\ell_i^S + (k/n)(\ell_i^{S+} - \ell_i^S)]$. This expression averages the realized losses over all uniform subsets of size k.

Appendix H. Empirical Uncertainty and Additional Results

Prediction intervals use the preceding validation year's 90th percentile of absolute errors for each procedure and reporting scenario. ARM endpoints are clipped to the feasible

TABLE A2. Paired squared-error differences and interval calculations

Target	Difference	Bank 95%	Holder 95%	Two-way 95%
2016				
ARM	8.64	[5.84, 11.45]	[5.84, 11.45]	[-2.49, 19.78]
Income	0.17	[0.10, 0.24]	[0.10, 0.24]	[-0.01, 0.35]
2022–2025, observed odd quarters				
ARM	21.04	[14.28, 27.80]	[14.26, 27.82]	[11.57, 30.51]
Income	0.08	[-0.23, 0.39]	[-0.24, 0.40]	[-1.25, 1.41]
2016–2025, observed odd quarters				
ARM	18.90	[15.83, 21.97]	[15.66, 22.14]	[15.81, 21.99]
Income	0.06	[-0.06, 0.19]	[-0.06, 0.19]	[-0.25, 0.38]

Differences equal semiannual loss minus available-report loss for the ridge procedure. Units are squared percentage points for ARM and squared basis points for income. Bank and holder intervals condition on the observed dates. Two-way intervals add target-quarter clustering; the smaller cluster count determines the critical value. The 2016 ARM calculation uses two target-quarter clusters.

share range. Reported coverage is the fraction of subsequent realizations inside those intervals. The serially dependent panel and annual distribution changes make this an empirical calibration exercise. Interval radii are empirical residual quantiles of fitted procedures. The calculations condition on estimated parameters.

Sampling summaries use paired squared-error differences. The bank-cluster variance sums squared bank totals of centered differences, with the cluster-count correction. Holding-company sensitivity groups certificates using the transitive union of their observed historical holding-company memberships. These groups enter uncertainty calculations only. The two-way calculation adds bank and target-quarter cluster components and subtracts their intersection component, following Cameron et al. (2011). Its t critical value uses the smaller cluster count minus one. Negative two-way variance estimates are omitted from interval reporting. In 2016, the two ARM target dates and four income target dates constrain this approximation. Year-specific differences and the count of target dates accompany pooled results.

Appendix I. Predictive Diagnostics

For test year Y, tuning and calibration use target year $Y - 1$, and fitting labels end by the first origin of Y. The specifications below examine predictive variance, bounded-share transformations, local scale, and repeated reported values under this chronological rule.

The expanded grid uses $\lambda_0 \in \{4, 16, 64, 256, 1024\}$, $a_0 \in \{1.1, 1.5, 2, 3, 6, 12\}$, and $q_0 \in \{0.0001, 0.001, 0.01, 0.025, 0.0625, 0.125, 0.25, 1, 4, 16, 64\}$. The criterion averages Student-t negative log density over the available and semiannual histories. Predictive second mo-

TABLE A3. Empirical coverage of 90 percent forecast intervals

Procedure	ARM semi.	ARM avail.	Income semi.	Income avail.
2016				
Ridge	89.5	89.2	47.7	47.7
Persistence	90.6	90.2	91.3	91.3
Gaussian filter	88.8	90.3	28.3	28.3
2022–2025, observed odd quarters				
Ridge	92.6	96.1	92.5	92.8
Persistence	92.6	96.1	94.4	94.4
Gaussian filter	90.6	96.5	92.6	92.7

Entries are percentages of subsequent outcomes within intervals calibrated on the preceding validation year. ARM intervals respect the share bounds. Coverage is an evaluated property of these intervals.

ments exist because $a_0 > 1$. The complete profiles include constant-history and varying-history subsets, with the latter diagnostic computed on the preceding-year validation sample. A constant history followed by the same value receives increasing log density as the scale approaches zero. Constant histories followed by a change generate the opposite pressure. Their relative frequencies help determine the observed boundary selection. Appendix I.3 estimates a likelihood for discrete persistence and continuous changes at the reporting resolution.

The variance-floor fit solves $\min_{\alpha,\beta\geq 0}\sum_i\{(Y_i-\widehat{m}_i)^2-\alpha-\beta V_i\}^2$ on the preceding year's available and semiannual histories. Test-year means use the baseline posterior, and predictive variance is $V_i^c=\alpha+\beta V_i$. A Student-t_5 distribution centered at that mean is scaled to variance V_i^c. The acquisition proxy is $g_i V_i^c/V_i$, evaluated on the semiannual archive. This calibration supplies a predictive variance and a ranking rule. Its coefficients are summaries of prior forecast errors.

The bounded specification transforms a share y in percentage points to $z=\log\{(y+\epsilon)/(100-y+\epsilon)\}$ with $\epsilon=0.1$. It fits the same conjugate increment model on z, using the displayed λ_0, a_0 grids and $q_0\in\{10^{-6},10^{-5},10^{-4},10^{-3},10^{-2},0.1,1,4\}$. The inverse map is $\max\{0,\min\{100,(100+2\epsilon)\,\mathrm{logit}^{-1}(z)-\epsilon\}\}$. Predictive moments use 513 Gauss–Legendre nodes in probability space. Comparisons with 1,025 nodes give maximum differences below 0.000063 percentage points in means and 0.019 squared points in variances. Its acquisition proxy multiplies the latent Gaussian gain by the squared derivative of the smooth inverse map at the latent predictive mean. This is a delta-method ranking, evaluated by realized forecast loss.

The local-scale benchmark centers its prediction on the last retained share. For retained increments Δy_r over gaps d_r, let $w_r=2^{-\ell_r/L}$, where ℓ_r is the age of the newer endpoint relative to the latest retained record. Its variance scale is $\widetilde{q}=(\nu\bar{q}+\sum_r w_r\Delta y_r^2/d_r)/(\nu+$

TABLE A4. Predictive calibration under alternative specifications

Specification	RMSE	Mean variance	Coverage (%)	Q1 variance	Q1 MSE
Baseline posterior	7.804	45.18	85.6	0.167	33.46
Expanded prior grid	7.808	60.34	87.0	0.009	33.46
Variance floor	7.804	67.12	96.4	45.430	33.46
Bounded transformation	7.909	76.89	87.1	0.032	33.45
Local scale, t_5	7.714	60.31	93.7	8.433	33.46

The 101,921 current-report origins are shared across specifications. Q1 fixes the lowest within-origin variance quintile under the baseline posterior (20,360 observations). Means and variances use percentage-point units; coverage concerns 90 percent predictive intervals. Each annual fit uses the preceding year. The variance floor calibrates squared errors; the local-scale benchmark uses a Student-t distribution with five degrees of freedom. These distributions are evaluated as predictive specifications.

TABLE A5. Reported-value resolution and the mixed predictive state

Band (bp)	Observed share	Predicted share	AUC	MSE/V	Both balances equal
0	19.50	19.46	0.996	0.97	0.06
1	20.61	20.56	0.989	0.97	0.06
5	23.27	23.20	0.968	0.97	0.06
10	25.63	25.45	0.950	0.97	0.06

Observed and predicted shares are percentages in the 2017–2025 out-of-time sample. A band of one basis point equals 0.01 percentage point of mortgage composition. Each row redefines the low-change component, refits the expanding-window classifier, and recomputes mixture variance. The last column reports exact equality of the ARM balance and mortgage-total balance between the latest retained report and the target report; it is repeated across band rows because its definition uses raw balances.

$\sum_r w_r$). The pooled $\overline{q}$ averages historical squared increments in preceding-year available and semiannual windows. The half-life $L \in \{1, 2, 4, 8\}$ and shrinkage weight $\nu \in \{1, 4, 16\}$ minimize preceding-year Student-t_5 negative log density. Predictive variance is $(d + h)\widetilde{q}$, and the acquisition proxy is $d\widetilde{q}$. This is a heavy-tail local-scale predictive specification. A latent stochastic-volatility transition law would introduce additional parameters.

Gradient boosting uses a HistGradientBoostingRegressor with squared loss, 120 iterations, learning rate 0.05, 15 maximum leaves, minimum leaf size 100, and L_2 penalty 1. Early stopping is disabled. Training stacks the available and semiannual histories with weight one half each and uses the current-bank and archive features. Training labels end in year $Y - 1$. Native missing-value handling is used, and share forecasts are clipped to $[0, 100]$. The hyperparameters are fixed across years. Persistence uses the last available report. For observed March/September origins, acquisition supplies the same latest record as available access.

I.1. The Common Mixed-State Acquisition Score

The score uses the estimated state for the intermediate report, semiannual archive features, and common bank controls available at the origin. Its label is the equal-weight mean of paired squared-error reductions under ridge, persistence, and gradient boosting in (35). Candidate boosting regressions have configurations (leaves, minimum leaf, L_2) equal to $(7, 250, 10)$, $(15, 200, 20)$, $(15, 250, 40)$, $(7, 500, 40)$, and $(31, 100, 10)$. The learning rate is 0.04, each fit has 220 iterations, and labels are winsorized at their 1st and 99th training percentiles.

The 2018 architecture is selected by five-fold bank cross-validation within target year 2017. For evaluation year $Y \geq 2019$, each candidate is fitted through $Y - 2$ and its acquisition value is measured in $Y - 1$; the winner is refitted through $Y - 1$ and applied in Y. Ties use the listed order. Each annual score is fixed before its evaluation outcomes. At budget b, the rule selects $\lfloor bn_t \rfloor$ reports within origin t, with certificate number breaking score ties.

For each budget, the paired value contrast uses the same bank, realized intermediate report, and later outcome under both allocation rules. Let $\widehat{\Sigma}_B$, $\widehat{\Sigma}_T$, and $\widehat{\Sigma}_{B\cap T}$ denote cluster-score covariance matrices by bank, target quarter, and their intersection. The budget-family covariance is $\widehat{\Sigma} = \widehat{\Sigma}_B + \widehat{\Sigma}_T - \widehat{\Sigma}_{B\cap T}$. A Gaussian maximum statistic over the 10, 25, and 50 percent budgets supplies simultaneous intervals, with its critical value multiplied by the ratio of the t_{15} and standard-normal 97.5 percent quantiles. Sixteen leave-one-target-quarter estimates measure time support. The main table compares the sequential score with uniform expected allocation and innovation-scale ranking. The three-budget grid was declared after the earlier 25 percent result; simultaneous inference covers emphasis within the displayed grid.

I.2. The Limit of the Continuous-Density Score

Let $U_{is} = \sum_r (\Delta y_r)^2 / d_r$ for the retained increments of scenario s. For a validation observation i whose retained increments are all zero, write $a_i = a_0 + (n_i - 1)/2$. Its posterior mean is the last retained share and its predictive scale squared is proportional to q_0. If the outcome repeats that share, its negative log density has coefficient $1/2$ on $\log q_0$. If the outcome changes, the Student-t tail gives coefficient $-a_i$. Histories with a positive squared increment have a positive limiting posterior scale for fixed $\lambda_0 > 0$. Thus, averaging equally over the two access scenarios s and N_Y validation rows,

(A11)
$$c_Y = \frac{1}{2N_Y} \sum_{s,i:\, U_{is}=0} \left\{ \frac{1}{2}\mathbf{1}(Y_i = y_{is,\mathrm{last}}) - a_{is}\mathbf{1}(Y_i \neq y_{is,\mathrm{last}}) \right\}, \qquad \mathrm{NLL}(q_0) = c_Y \log q_0 + O(1).$$

At $a_0 = 1.1$, the ten annual coefficients range from 0.03169 to 0.06634; at $a_0 = 3$, they range from 0.01309 to 0.05379. Their signs imply an unbounded criterion as q_0 approaches zero.

This result explains the boundary behavior of the continuous-density tuning criterion. Appendix I.3 uses a mixed likelihood for the predictive state. The fixed-grid score remains a benchmark input to the common acquisition rule.

I.3. A Point-Mass and Change Likelihood

Let $R_{it} = 1$ when the target share equals the latest share in the semiannual archive at the numerical precision of the constructed reported percentage, and let $p_{it} = \Pr(R_{it} = 1 \mid X_{it})$. Conditional on a change, the target has a Student-t_5 density $f_5(y; m_{it}, v_{it})$ centered on the expanded-grid predictive mean. The likelihood contribution relative to the sum of a point mass at the latest share and Lebesgue measure is

$$\mathcal{L}_{it} = p_{it}^{R_{it}}\{(1 - p_{it}) f_5(Y_{it}; m_{it}, v_{it})\}^{1-R_{it}}. \tag{A12}$$

The probability model is a histogram gradient-boosting classifier with learning rate 0.05, 120 iterations, at most 15 leaves, minimum leaf size 150, and L_2 penalty one. Its inputs are the log baseline variance, absolute drift, log innovation estimate, record count, age, span, latest share, predicted change, log assets, calendar quarter, Call Report form, forecast horizon, origin date, and mortgage-to-asset ratio. The feature set excludes the current ARM share under semiannual access. For target year Y, all target years before Y form the training sample. The first prediction year is 2017.

The conditional change variance uses twenty cells defined by five-percent quantiles of the expanded-grid variance among training observations with $R_{it} = 0$. Within a cell, v_{it} equals the training mean of $(Y_{it} - m_{it})^2$. Test observations use the training cutoffs and cell variances. The mixture mean and variance are

$$\mu_{it} = p_{it}\ell_{it} + (1 - p_{it})m_{it}, \qquad V_{it} = (1 - p_{it})v_{it} + p_{it}(1 - p_{it})(\ell_{it} - m_{it})^2, \tag{A13}$$

where ℓ_{it} is the latest retained share. Quantiles invert the continuous t_5 distribution on each side of its jump of size p_{it} at ℓ_{it}. The mixed log-score comparison in the text changes p_{it} to the expanding-window repeat frequency and keeps f_5 fixed, so both scores use the measure in (A12).

For the resolution sensitivity, $R_{it}(b) = 1\{|Y_{it} - \ell_{it}| \le b/100\}$ for $b \in \{0, 1, 5, 10\}$ basis points of composition. Each definition refits the annual classifier, re-estimates the conditional change-variance cells, and recomputes mixture moments. The accounting comparison uses the ARM numerator and first-lien mortgage denominator for the target report and the latest retained report. Exact share repetition occurs in 19.76 percent of the full evaluated sample; the ARM numerator is equal in 19.01 percent, the denominator is equal in 0.08 percent, and both are equal in 0.06 percent. At the intermediate report origin, the

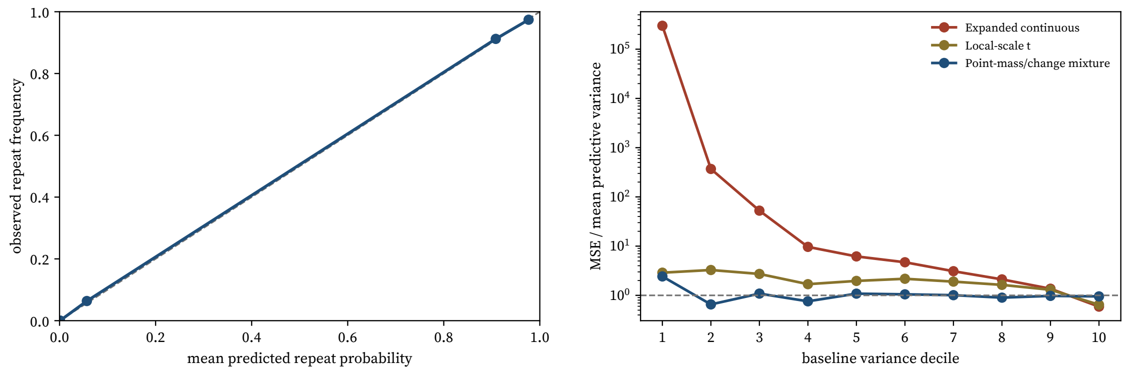

FIGURE A2. Repeat probabilities and predictive-variance calibration

The left panel groups the 2017–25 out-of-time observations into deciles of predicted repeat probability. The right panel groups them into deciles of the baseline semiannual variance and reports realized MSE divided by mean predictive variance. The horizontal line marks equality. Each annual fit uses earlier target years.

corresponding shares are 59.81, 59.75, 49.92, and 49.89 percent. These counts distinguish a repeated ratio from repeated accounting inputs.

Appendix J. Reform Cohorts and Eligibility

The cohort construction fixes domestic commercial-bank cohorts at June 2016 and June 2018, using bank classes N, NM, and SM and forms 041 or 051. The first cohort contains 4,654 banks below one billion dollars; the second contains 455 banks with assets from one billion to below five billion dollars. Subsequent presence is measured in the same bank-class/form sample. These are certificate-based cohort histories; merger exits and changes in bank classification alter the present denominator.

The 2019 eligibility rule uses preceding-June assets and also imposes conditions concerning foreign offices, advanced approaches, and institution classification, with supervisory reservation of authority (Board of Governors of the Federal Reserve System, Federal Deposit Insurance Corporation, and Office of the Comptroller of the Currency 2019). The asset/class/form criteria in this cohort construction are observable proxies. Voluntary adoption can depend on expected growth and the costs of maintaining an existing reporting system. The 2017 one-billion-dollar boundary also coincides with examination-cycle rules associated with the same asset boundary. A causal use of eligibility would require a first-stage analysis, exclusion restrictions addressing concurrent changes, and a treatment of selection into continued observation. The 2019 ten-percent window has zero observed 051 adoption on either side in September 2019 and March 2020.

TABLE A6. Income forecast information and the marginal ARM record

Information set	Semiannual RMSE	Available RMSE	MSE gain (%)	MSE reduction
2016: main sample; $N = 19,480$				
All controls	10.778	10.770	0.146	0.170
Earnings removed	16.024	15.975	0.600	1.539
Maturity removed	10.832	10.829	0.052	0.061
Both groups removed	14.922	14.889	0.444	0.988
2022–2025: observed March/September; $N = 7,085$				
All controls	11.069	11.066	0.065	0.080
Earnings removed	21.620	21.632	-0.106	-0.497
Maturity removed	11.042	11.037	0.089	0.108
Both groups removed	21.500	21.513	-0.118	-0.546
2016–2025: observed March/September; $N = 29,588$				
All controls	11.623	11.620	0.047	0.064
Earnings removed	18.473	18.456	0.182	0.622
Maturity removed	11.684	11.681	0.038	0.051
Both groups removed	18.179	18.165	0.149	0.494

The shared eight-scenario ridge procedure and rolling samples coincide across restrictions. Earnings comprise four current/lag/change/availability features. Maturity comprises six mortgage, six other-loan, eight deposit, and twelve securities bins. Remaining balance sheets, rates, and archive inputs stay available. RMSE is in basis points of origin assets; MSE reductions are in squared basis points. Positive gains favor the available-report archive.

Appendix K. Income Ablations

The same records can be evaluated against next-quarter net interest income. A shared ridge coefficient vector for each outcome takes current bank characteristics, maturity measurements, earnings, rates, and archive summaries as inputs. Coefficients are shared across eight access restrictions. For year y, fitting uses outcomes through $y-2$, penalty selection uses $y-1$, and refitting uses outcomes through $y-1$. Comparisons use identical bank-origin rows. Appendix G gives the procedure and its comparators.

In the 2016 schedule comparison, quarterly access reduces composition MSE by 17.02 percent and income MSE by 0.15 percent. Table A6 removes the earnings group, the maturity group, and both groups while preserving the fitting rule, remaining predictors, and sample. Earnings comprise current income, its lag, its change, and the lag-availability indicator. Maturity comprises six mortgage, six other-loan, eight deposit, and twelve securities measurements. These restrictions evaluate marginal ARM record value under different existing income information. Removing earnings raises available-archive income RMSE from 10.770 to 15.975 basis points in 2016, and from 11.066 to 21.632 in the observed March and September sample during 2022–2025. Its effect on the marginal ARM record varies across those periods. The income MSE gain rises from 0.146 to 0.600 percent in 2016 and changes from 0.065 to −0.106 percent in 2022–2025. In the pooled observed-origin sample, the increase in the paired ARM gain after removing earnings is 0.558 squared basis points, with two-way interval $[-0.347, 1.463]$.

Current earnings reduce overall income forecast error in these comparisons. The ARM contribution stays small and varies with the period and remaining predictors. This target dependence corresponds to the covariance term in Lemma 4; the reported ablations estimate forecast losses under the specified procedures. A policy-loss matrix would additionally require the exposure of policy actions to those income outcomes. The 2016 persistence income forecast has RMSE 5.761 basis points, compared with 10.770 for full-information ridge. Its 90 percent interval coverage is 91.3 percent; ridge coverage is 47.7 percent. Appendix H reports forecast errors and coverage for every procedure. The ablations describe a fitted income forecast's response to its information set, including its estimation error.

Appendix L. Reporting primitives and dynamic extensions

L.1. Operating Priorities

The static direction choices can be represented by value and exposure moments. Let $p, \zeta \in \mathbb{R}^d$ denote operating values and exposures, and let $M_r = E[p\zeta']$. An organization chooses V to maximize $E[p'V\Sigma W'\zeta] = \mathrm{tr}(B_r'V)$, where $B_r = M_r W\Sigma$. For nonsingular B_r, its choice is

$$\bar{V}_r = B_r(B_r'B_r)^{-1/2}. \tag{A14}$$

A statistical agency has monitoring values $C_r = E[\nu\nu']$ and chooses U to maximize $\mathrm{tr}(C_r U\Lambda U')$. Its columns align ordered eigenvectors of C_r with the ordered eigenvalues of Λ. These operating problems determine desired directions conditional on the regime. Their derivations are in Appendix A.5.

The dynamic model below uses quadratic costs of deviations from those desired directions and of changing an installed direction. The quadratic operating costs are maintained primitives on an angular coordinate. They can represent local approximations to the operating criteria above; the dynamic results use them as the specified objective. Both decision makers observe the priority path. Their objectives are separable conditional on that path. The later policymaker benefits from the observations, and its loss is assigned to the planner's objective with weight χ. This allocation of objectives specifies the source of the information externality.

L.2. One Report and Two Posterior States

The baseline state (x, J) records installation and precision when future target demand is known. A report can also reveal which policy problem will use the archive. Let $\theta \sim N(0, S)$ with $S = \mathrm{diag}(s_1, s_2)$, and let a latent target index satisfy $T = \alpha + \lambda\theta_1 + \zeta$, where $\zeta \sim N(0, \sigma_T^2)$ is independent and $\lambda > 0$. The target regime is $Z = 1\{T \geq 0\}$. Conditional on $Z = 1$, future policy demand assigns weight d_H to the first coordinate and d_L to the second; the weights are reversed under $Z = 0$, where $d_H > d_L \geq 0$.

The installed system first reports

$$Y = \theta_1 + \epsilon, \qquad \epsilon \sim N(0, r), \quad r > 0. \tag{A15}$$

Write $k = s_1/(s_1 + r)$, $p_1 = s_1 r/(s_1 + r)$, and $\sigma_*^2 = \sigma_T^2 + \lambda^2 p_1$. Gaussian conditioning gives

$$\begin{aligned} \mu(y) &= E(\theta \mid Y = y) = (ky, 0)', \qquad P = \mathrm{Var}(\theta \mid Y = y) = \mathrm{diag}(p_1, s_2), \\ \pi(y) &= \Pr(Z = 1 \mid Y = y) = \Phi\left(\frac{\alpha + \lambda k y}{\sigma_*}\right). \end{aligned} \tag{A16}$$

Thus the first report changes coefficient precision from S^{-1} to $J = P^{-1}$ and its realized content changes conditional target demand,

$$D\{\pi(y)\} = \mathrm{diag}\{d_L + (d_H - d_L)\pi(y),\ d_H - (d_H - d_L)\pi(y)\}. \tag{A17}$$

The next report has precision $\tau > 0$ and is placed at e_1 or e_2. A move from the installed coordinate e_1 to e_2 costs $c > 0$. Put $A(p) = \tau p^2/(1 + \tau p)$. Conditional on $Y = y$, the expected posterior policy loss after placing the next report at e_i is

$$R_i(y) = \frac{1}{2}\,\mathrm{tr}\left[D\{\pi(y)\}\left\{P - \frac{\tau P e_i e_i' P}{1 + \tau e_i' P e_i}\right\}\right]. \tag{A18}$$

THEOREM A1 (A report-content cutoff after a precision update). *Define $A_1 = A(p_1)$, $A_2 = A(s_2)$, $\Delta_d = d_H - d_L$, and*

$$\pi^\dagger = \frac{d_H A_2 - d_L A_1 - 2c/\chi}{\Delta_d (A_1 + A_2)}. \tag{A19}$$

For $\chi > 0$, the planner moves the next report to e_2 exactly when $\pi(Y) < \pi^\dagger$ and retains e_1 exactly when $\pi(Y) > \pi^\dagger$. If $\pi^\dagger \in (0, 1)$, the report rule has the realization cutoff

$$Y < y^\dagger, \qquad y^\dagger = \frac{\sigma_* \Phi^{-1}(\pi^\dagger) - \alpha}{\lambda k}. \tag{A20}$$

Let $v(y) = \min\{\chi R_1(y), c + \chi R_2(y)\}$. A plan that selects the second direction after observing the first report satisfies

$$E\{v(Y)\} \le \min\left[E\{\chi R_1(Y)\}, E\{c + \chi R_2(Y)\}\right]. \tag{A21}$$

The inequality is strict when Y reaches both action regions with positive probability.

The first observation supplies both arguments of the second reporting decision. Its measurement precision determines P, and its realization determines $D\{\pi(Y)\}$. The cutoff compares target-weighted variance reductions $d_1\{\pi(y)\}A_1$ and $d_2\{\pi(y)\}A_2$ with the movement cost. A larger inherited variance along a coordinate raises the variance reduction available from measuring that coordinate. A realization associated with target regime one raises the weight on the first coordinate and shifts the second report toward e_1. The content state extends the planner recursion to (x, μ, P, π). After a report realization, Bayes' rule maps this state into (μ', P', π'), and future directions can depend on all four entries. Equation (A21) prices the option to choose the next report after observing the first one. The timing uses the report once: it enters the coefficient posterior used in policy loss and the target posterior used in policy demand.

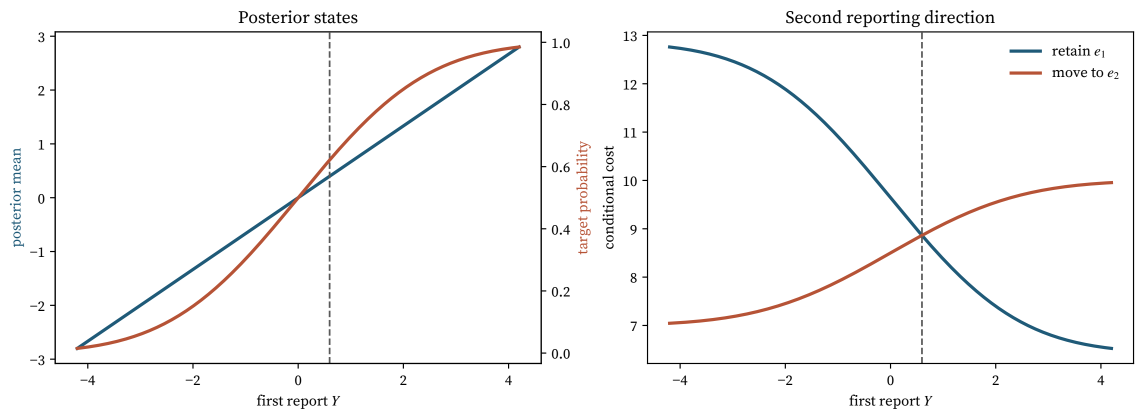

FIGURE A3. One report updates the archive and target demand

The left panel maps the first report realization into the posterior coefficient mean and the probability of target regime one. Coefficient precision changes from S^{-1} to P^{-1} for every realization. The right panel plots conditional costs for the two possible second reports; their crossing gives $y^{\dagger}$. Parameters are $S = \operatorname{diag}(2, 1)$, $r = 1$, $\alpha = 0$, $\lambda = 1$, $\sigma_T^2 = 1$, $\tau = 2$, $d_H = 4$, $d_L = 1$, $\chi = 6$, and $c = 1$.

L.3. Reporting Paths and the Archive They Create

Consider H reporting dates, a stable coefficient vector, a constant zero priority, $\delta_0 = 0$, $D_t = I_2$, and rank-one batches with precision $\tau e(\delta_t)e(\delta_t)'$, where $\tau > 0$. Let the initial archive be $J_0 = \operatorname{diag}(a, b)$ with $a \geq b > 0$. The objective in (21) becomes

(A22)
$$F_H(\delta) = \frac{1}{4}\sum_{t=1}^{H}\eta^{t-1}\{\rho\delta_t^2 + \kappa(\delta_t - \delta_{t-1})^2\} + \frac{\chi}{2}\sum_{t=1}^{H}\eta^{t-1}\operatorname{tr}(J_t^{-1}), \qquad J_t = J_0 + \tau\sum_{s=1}^{t} e(\delta_s)e(\delta_s)'.$$

Here $\rho, \kappa > 0$ and $0 < \eta < 1$. The report at each date changes the archive used for every remaining policy decision. Writing $S_t = a + b + t\tau$ and $K_t = b(a + t\tau)$, the determinant is

$$\det J_t = K_t + \tau(a - b)\sum_{s=1}^{t}\sin^2\delta_s + \tau^2\sum_{1\leq r<s\leq t}\sin^2(\delta_r - \delta_s). \tag{A23}$$

The last sum records information produced by differences between reporting dates. It enters even when the initial archive has $a = b$.

Define symmetric $H \times H$ matrices A_H and B_H through their quadratic forms, with $z_0 = 0$:

$$\begin{aligned} z'A_H z &= \sum_{t=1}^{H}\eta^{t-1}\{\rho z_t^2 + \kappa(z_t - z_{t-1})^2\}, \\ z'B_H z &= \sum_{t=1}^{H}\frac{\eta^{t-1}S_t}{K_t^2}\left\{\tau(a-b)\sum_{s=1}^{t} z_s^2 + \tau^2\sum_{1\leq r<s\leq t}(z_r - z_s)^2\right\}. \end{aligned} \tag{A24}$$

THEOREM A2 (A threshold for a sequence of reporting choices). *The objective* (A22) *has a minimizer. If $B_H \neq 0$, define*

$$\chi_H^* = \frac{1}{2\lambda_{\max}(A_H^{-1/2} B_H A_H^{-1/2})}. \tag{A25}$$

For $0 \leq \chi \leq \chi_H^*$, the unique minimizing path is $\delta_t = 0$ at every date. For $\chi > \chi_H^*$, every minimizing path has a nonzero angle. If $B_H = 0$, the zero path is uniquely optimal at every finite χ. With $a = b$, $B_1 = 0$ and $B_H \neq 0$ for $H \geq 2$. With $H = 1$ and $a > b$, (A25) *equals the threshold in Theorem 2 for $\lambda_L = 0$.*

The threshold compares the operating and adjustment costs of an entire path with its effects on all subsequent archives. The proof uses (A23) to bound the objective over every path; the generalized eigenvalue then determines where the zero path first loses optimality. The two-date case also yields every minimizing path.

COROLLARY A1 (Reporting costs and the path threshold). *Fix H, a, b, τ, η with $B_H \neq 0$. The threshold χ_H^* is strictly increasing in each of ρ and κ. Scaling both costs by $c > 0$ scales the threshold by c.*

These comparisons keep information production fixed. The discounted difference operator with $z_0 = 0$ and the operating-cost matrix are positive definite, so each cost increase raises the Rayleigh quotient defining the threshold. Appendix A.6 gives the argument. For two dates, Corollary A2 also signs the discount-factor comparison. Figure A1 displays longer-horizon calculations for a stated parameter grid.

COROLLARY A2 (Two dates with equal initial precision). *Let $H = 2$ and $a = b = j > 0$. Define $h = \rho + \kappa$, $K = j(j + 2\tau)$, $S = 2j + 2\tau$, and*

$$c = \kappa + \frac{\rho h}{h + \eta\rho}, \qquad \chi_2^* = \frac{cK^2}{2S\tau^2}. \tag{A26}$$

For $\chi \leq \chi_2^*$, both reporting angles are zero. For $\chi > \chi_2^*$, there are two minimizing paths, related by reflection. The path with positive separation $z = \delta_2 - \delta_1$ satisfies

$$\delta_1 = -\frac{\eta\rho}{h + \eta\rho} z, \qquad \delta_2 = \frac{h}{h + \eta\rho} z, \qquad cz = \frac{\chi S \tau^2 \sin 2z}{(K + \tau^2 \sin^2 z)^2}, \qquad 0 < z < \frac{\pi}{2}. \tag{A27}$$

There is one such positive z. It increases with χ/c, tends to zero as χ decreases to χ_2^*, and tends to $\pi/2$ as χ increases without bound. The threshold increases with κ and decreases with η.

The first report by itself has the same expected policy loss at every angle because the initial archive and demand are isotropic. Its chosen angle anticipates the second report. Above the threshold, the institution first moves to one side of its operating priority and then moves to the other side. Its first observation creates an archive with unequal precision, and its second observation uses the less measured direction. The operating priority and policy demand stay constant throughout this sequence. The result characterizes two dates; equation (A25) gives the onset of reorientation for any finite horizon.

To examine the role of the isotropic benchmark, let $F_2(\delta;\zeta)$ denote the two-date objective with $J_0 = \mathrm{diag}(a, b)$, $D = \mathrm{diag}(d_1, d_2)$, first- and second-report precisions (τ_1, τ_2), and the operating parameters in (A22). The vector ζ collects these positive quantities. Let ζ_0 set $a = b = j$, $d_1 = d_2 = 1$, and $\tau_1 = \tau_2 = \tau$ while retaining the other parameters in Corollary A2.

THEOREM A3 (Straddling under nearby precision and target demand). *Fix $\chi > \chi_2^*$ at ζ_0, and let $\delta^0 = (\delta_1^0, \delta_2^0)$ be the minimizing path with $\delta_1^0 < 0 < \delta_2^0$. There are a neighborhood $\mathcal{U}$ of ζ_0, two continuously differentiable functions $\delta^+(\zeta)$ and $\delta^-(\zeta)$, and a finite constant C such that every global minimizer of $F_2(\cdot;\zeta)$ for $\zeta \in \mathcal{U}$ equals one of these two branches and*

$$\min_{q\in\{\delta^0,-\delta^0\}} \|\delta^s(\zeta) - q\| \le C\|\zeta - \zeta_0\|, \qquad s \in \{+,-\}. \tag{A28}$$

Each branch places its two reports on opposite sides of the zero operating priority: $\delta_1^s(\zeta)\delta_2^s(\zeta) < 0$.

The theorem uses the positive curvature at the two benchmark minima and the objective gap separating them from the remaining paths. Precision and target-demand asymmetry can therefore change the two angles while preserving their ordering around the operating priority. The neighborhood depends on the benchmark distance above χ_2^* and the angular margins in (A27).

Appendix M. Archive-state estimation and predictive calibration

M.1. A Posterior State from Retained Reports

The empirical state describes a changing mortgage share under a Gaussian-increment approximation. Over a retained window, let the increment between reports separated by d quarters satisfy $y_{i,s+d} - y_{is} = db_i + \epsilon_{is,d}$, where $\epsilon_{is,d} \mid q_i \sim N(0, dq_i)$ and disjoint increments are independent. The drift prior is $b_i \mid q_i \sim N(0, q_i/\lambda_0)$, and q_i has an inverse-gamma prior. The oldest retained report anchors the level. Reports observe composition exactly; uncertainty about its subsequent movement depends on the retained increments. The two-coordinate state consists of current composition and its drift, with innovation variance learned from the same history.

For archive H_{it}, let d_{it} be the age of its latest report, $\widehat{b}_{it} = E[b_i \mid H_{it}]$, $\widehat{q}_{it} = E[q_i \mid H_{it}]$, and $\lambda_{it} = \lambda_0 + \mathrm{span}(H_{it})$. The predictive mean, variance, and inverse posterior drift variance

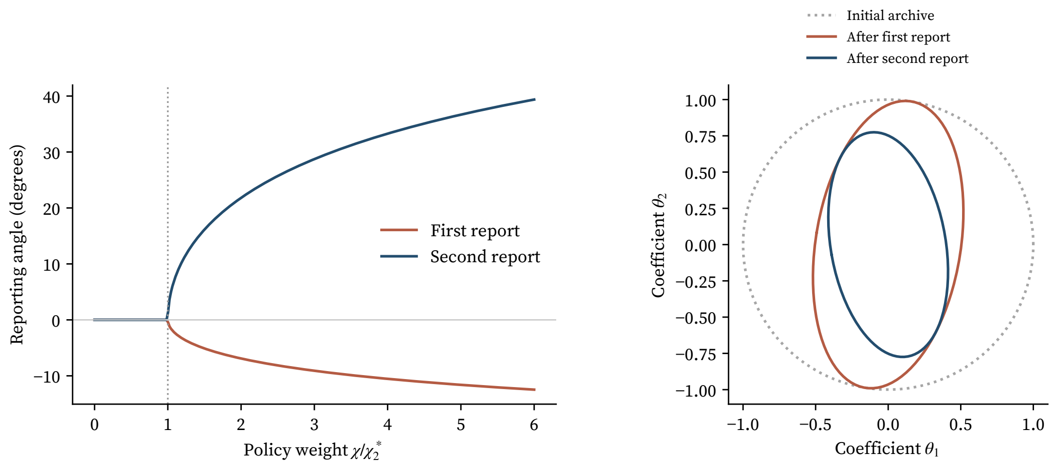

FIGURE A4. Jointly chosen reports under a constant priority

The calculations use $j = 1$, $\tau = 3$, $\rho = 1$, $\kappa = 2$, and $\eta = 0.95$. The left panel plots the two minimizing angles on the branch with positive separation against χ/χ_2^*. The right panel traces the posterior covariance ellipses after the first and second reports at $\chi = 3\chi_2^*$. These are model calculations. The reflection of the displayed path has the same cost.

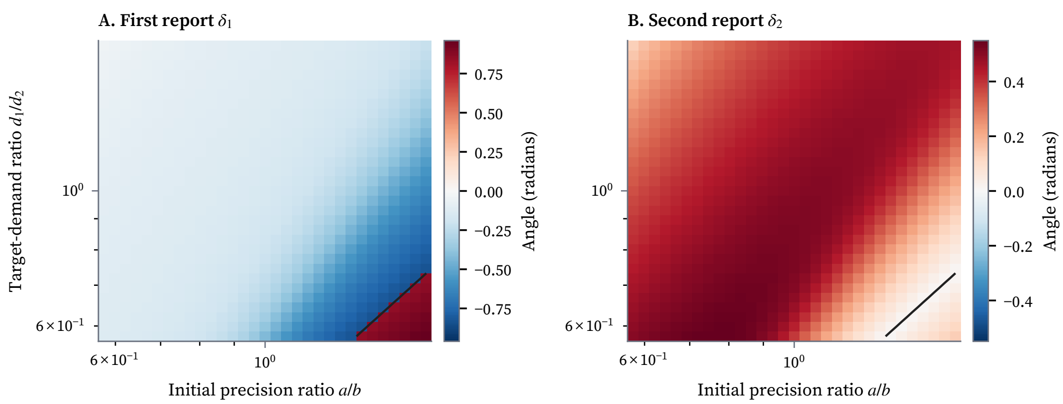

FIGURE A5. Two-date reporting with asymmetric precision and demand

The calculations set $j = 1$, $\tau_1 = \tau_2 = 3$, $\rho = 1$, $\kappa = 2$, $\eta = 0.95$, and $\chi = 3\chi_2^*$. The horizontal axis changes initial precision through a/b while preserving $ab = 1$; the vertical axis changes policy demand through d_1/d_2 while preserving $d_1 d_2 = 1$. Each grid point uses 25 starting values for the two-angle global search. Both reports straddle zero throughout the rectangle $a/b, d_1/d_2 \in [0.8, 1.25]$.

are

$$\text{(A29)}\qquad \begin{aligned} \widehat{m}_{it}(h) &= y_{i,\text{last}} + (d_{it} + h)\widehat{b}_{it}, \\ V_{it}(h) &= \widehat{q}_{it}\left\{ d_{it} + h + \frac{(d_{it}+h)^2}{\lambda_{it}} \right\}, \qquad j_{it}^{b} = \frac{\lambda_{it}}{\widehat{q}_{it}}. \end{aligned}$$

At a current report, $d_{it} = 0$ and the observed level coincides across retained histories. Those histories can still imply different drift uncertainty and innovation variance. With q_i

TABLE A7. Retained histories with the same current report

Period	Records	Origins	RMSE	Mean V	MSE gain (%)
2016	4	18,070	6.381	21.04	0.00
2016	8	18,070	6.361	31.67	0.62
2016	12	18,070	6.369	37.02	0.36
2022–2025	4	33,482	7.608	27.35	0.00
2022–2025	8	33,482	7.608	40.94	0.01
2022–2025	12	33,482	7.611	48.26	-0.07
2016–2025	4	101,921	7.821	27.07	0.00
2016–2025	8	101,921	7.807	39.63	0.35
2016–2025	12	101,921	7.804	45.18	0.43

Every row within a period uses the same bank-origin observations and current ARM report. RMSE is in percentage points and variance in squared percentage points. MSE gains use four records as the reference. Posterior innovation variance is learned separately from each retained history.

fixed, a longer span reduces the drift contribution to future variance. When q_i is learned, older records also revise the estimated amount of variation. The full posterior includes that variance state. Appendix D derives its Student-t predictive distribution.

Prior tuning parameters are selected by predictive log score on the preceding year's outcomes and held fixed for the target year. Bank-specific posteriors use retained records dated through the origin. Appendix D.3 gives the annual selection rule and predictive diagnostics.

M.2. Holding the Current Report Fixed

There are 101,921 bank-origin observations with a current ARM report and twelve available records. Supplying four, eight, or twelve records keeps the current report and all other current bank records identical within each comparison. Table A7 reports future forecast losses and posterior variances. Over 2016–2025, increasing the archive from four to twelve reports changes RMSE from 7.821 to 7.804 percentage points, a 0.43 percent MSE reduction. The two-way bank and target-date interval for the paired MSE reduction is $[-0.12, 0.65]$ squared percentage points. Average predictive variance rises as longer histories revise estimated innovation variance. The uncertainty state is evaluated against subsequent errors. Within each origin, banks are grouped by current composition decile and asset quintile. Removing the means of these joint cells, the slope of realized squared error on $V_{it}(h_t)$ is 0.381, with a two-way interval of $[0.297, 0.464]$. Table A8 reports the first year, 2022–2025, and the pooled sample. The association uses future errors and an archive state calculated at the origin. The paired retention comparison fixes the current observation exactly; the cross-bank calibration uses the stated bins.

TABLE A8. Archive states and subsequent forecast losses

Prediction	Period	Origins	Slope	Bank 95% CI	Two-way 95% CI
Uncertainty	2016	18,070	0.334	[0.161, 0.507]	[−0.182, 0.849]
Uncertainty	2022–2025	33,482	0.448	[0.315, 0.581]	[0.288, 0.608]
Uncertainty	2016–2025	101,921	0.381	[0.304, 0.457]	[0.297, 0.464]
Report gain	2016	9,799	0.402	[0.136, 0.667]	—
Report gain	2022–2025	7,085	0.338	[0.060, 0.617]	[−0.075, 0.751]
Report gain	2016–2025	29,588	0.374	[0.222, 0.526]	[0.276, 0.472]

Uncertainty regresses subsequent squared error on predicted variance after removing joint origin, current-composition-decile, and asset-quintile cell means. Report gain regresses realized paired MSE reduction on its ex ante prediction, with cells using latest semiannual composition. Two-way intervals cluster by bank and ARM target date. The 2016 sample has two target dates. A dash records a negative estimated two-way variance. Slopes condition on fitted prior parameters.

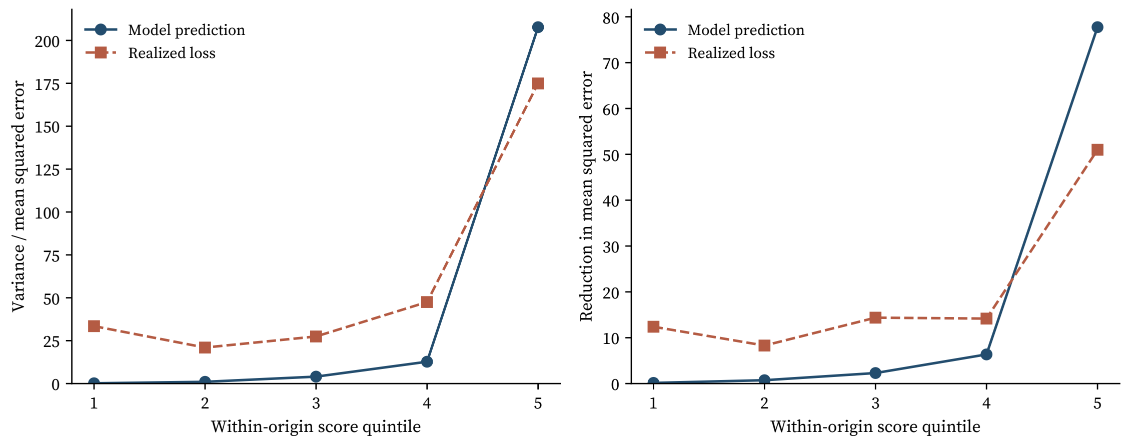

FIGURE A6. Archive uncertainty and the value of an additional record

The left panel uses 101,921 current-report origins with twelve retained observations during 2016–2025. Quintiles use model predictive variance within each origin; the lines compare average variance with subsequent squared error. The right panel uses 29,588 observed March and September origins. Quintiles use expected MSE reduction before the current report is supplied; the lines compare it with realized paired loss reduction. Units are squared percentage points. Forecasts and scores use the annual out-of-time fitting rule.

Figure A6 shows the scale of the calibration error. In the lowest within-origin variance quintile, average predicted variance is 0.17 and realized MSE is 33.46 squared percentage points. In the highest quintile, the corresponding values are 207.72 and 174.93. Large composition changes occur among banks with quiet retained histories. These outcomes qualify the absolute uncertainty estimates even where the archive state predicts variation in subsequent losses.

TABLE A9. Predictive calibration with repeated reported values

Model	RMSE	MSE/V	Cov.	Quiet MSE/V	Quiet cov.
Expanded continuous	8.692	1.09	87.0	748.90	91.3
Local-scale t	8.592	1.03	94.0	3.07	98.3
Point-mass/change	8.664	0.97	95.5	1.04	97.4

Coverage columns report percent coverage of 90 percent predictive intervals; V denotes mean predictive variance. Quiet histories are the lowest fifth of baseline semiannual predictive variance in the 2017–25 out-of-time sample. The mixture assigns an atom to the latest retained value and a Student t_5 distribution to changes. Each target year uses earlier target years to estimate repeat probabilities and twenty conditional-variance cells.

M.3. Prior Profiles and Predictive Calibration

Extending the prior grid toward smaller innovation scales places the continuous-density criterion at its lower boundary in every validation year. Conditional profiles separate constant retained histories from histories with changes, and Appendix I.2 gives the boundary calculation. A point-mass/change likelihood then assigns probability to repeated reported values and a continuous distribution to changes.

Across alternative variance specifications, the quiet-history group remains the main source of absolute calibration error. The point-mass/change model addresses the repeated-value component directly and is evaluated out of time. The online appendix reports the alternative variance fits and interval coverage. The sequential mixture uses 2016 as its first training year and predicts each target year from 2017 through 2025 using earlier target years. Exact repetition of the latest retained reported share occurs in 19.50 percent of the 145,569 out-of-time observations; the mean predicted probability is 19.46 percent. Annual repeat classifiers have AUC between 0.993 and 0.997. Relative to an expanding-window constant repeat probability with the same change distribution, the fitted probability lowers mixed negative log score by 0.446 per observation. Table A9 reports predictive calibration. The mixture's MSE-to-variance ratio is 0.97 overall and 1.04 in the quiet-history fifth, compared with 3.07 for the local-scale benchmark. Its central intervals cover 95.5 percent overall and 97.4 percent in that fifth; probability mass at a reported value can make central-interval coverage exceed its nominal level. Appendix I.3 gives the likelihood, features, chronology, and decile diagnostics. Across 0, 1, 5, and 10 basis-point resolution bands, the observed low-change share rises from 19.50 to 25.63 percent, the mean fitted probability follows within 0.18 percentage point, and the predictive MSE-to-variance ratio stays between 0.970 and 0.972. Exact equality of the target ARM balance occurs in 19.01 percent of origins, while equality of both the ARM balance and mortgage-total balance occurs in 0.06 percent. The repeated-share component therefore largely reflects an unchanged ARM balance with a changing mortgage denominator. Appendix I gives the parameter grids, transformations, moments, and fitting rules. The next calculation evaluates acquisition rankings using separate predictors and realized losses.

References

Board of Governors of the Federal Reserve System, Federal Deposit Insurance Corporation, and Office of the Comptroller of the Currency. Reduced reporting for covered depository institutions, 2019. URL `https://www.federalreserve.gov/newsevents/pressreleases/files/bcreg20190617a1.pdf`. Final rule released June 17, 2019; reporting changes effective September 30, 2019.

A. Colin Cameron, Jonah B. Gelbach, and Douglas L. Miller. Robust inference with multiway clustering. *Journal of Business & Economic Statistics*, 29(2):238–249, 2011. doi: 10.1198/jbes.2010.07136.