Economica Obscura

Institutions, Archives, and Policy Learning

Chae-Yeon Xon
Yonsei University

September 2026

Abstract

Reporting institutions bear current collection and adjustment costs, while later users benefit from the records they leave. This paper models that information externality. Reporting directions adapt to operating priorities and generate an archive of state information. In a Gaussian benchmark, specialization can leave finite precision in a neglected direction even as precision grows in the favored direction. A planner values future policy losses as well as institutional costs; inherited precision determines when redirecting the next report is worthwhile. An archive-value payment implements the planner comparison under the stated commitment and target assumptions. A U.S. Call Report exercise evaluates the predictive contribution of an intermediate report before its contents are supplied. Prior-year acquisition scores improve loss over uniform selection, while innovation scale explains much of the gain. Forecast-loss evidence, whole-form burden measures, and the unmeasured monetary value of information enter distinct parts of the institutional interpretation.

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.

1. Introduction

A report collected for today's mandate becomes part of tomorrow's information set. A bank supervisor chooses which fields to require and how often to collect them; later users inherit those records when forecasting an exposure or evaluating another policy. The institution bearing the reporting cost and the user receiving the future information benefit can face different objectives. The resulting externality concerns the content of the archive as well as its size.

A two-direction reporting problem isolates the incentives. One direction serves the current operating priority; the other can help a future user whose target differs. Redirecting a report costs the current institution, while its information remains in the archive. In the model, departures from operating priorities and changes to installed systems have separate costs. Successive reports add precision about the economic state. The planner combines those institutional costs with the later user's target-weighted loss, making explicit whose benefit is left out of the decentralized choice.

Under a fixed priority, the reporting direction specializes gradually. The specialization theorem shows that, in the stable-coefficient rank-one Gaussian benchmark, precision grows along the limiting direction while remaining finite in the neglected direction; posterior policy loss then converges to an explicit positive limit whenever the future target loads on that direction. Complete learning holds exactly when the cumulative number of supplemental reports in the neglected direction diverges, so their reporting frequency may vanish. A planner can slow specialization, and an archive-value payment implements the planner path under the stated commitment and target assumptions. Inherited precision also creates an explicit threshold at which the next report turns away from the current priority. State persistence determines how much an older record contributes to that comparison.

U.S. Call Reports supply an empirical setting for this margin. Changes in reporting schedules and form adoption describe installed reporting choices. Retained semiannual mortgage reports form a baseline archive; the exercise asks which intervening March or September report to acquire before seeing its contents.[1] Candidate scores use the available history and earlier target years. Evaluation compares paired forecast losses after the additional report is supplied, holding the stated report budgets and prediction procedures fixed.

At 10, 25, and 50 percent budgets over 2018–2025, the sequential score improves equal-weight forecast loss over uniform acquisition, with positive simultaneous lower endpoints under the reported checks. An innovation-scale rule produces similar allocations, and its difference from the richer score is estimated less precisely. The evidence therefore

[1] The acquisition unit is one dated bank report. Its score is constructed from information available before the report's contents are revealed.

identifies useful archive information for ranking reports, while also showing the strength of a simpler benchmark.

The theory and the evidence enter the institutional conclusion at different margins. The model describes optimal reporting paths and an archive-value payment. The empirical exercise measures forecast-loss gains and report counts. Whole-form preparation hours provide a cost scale; item-level costs, policy actions, and the conversion of forecast loss to money remain additional inputs to welfare. The paper reports this mapping explicitly.

Measurement choices respond to objectives and institutional incentives (Almås et al. 2024; Veldkamp and Chung 2024; Liang and Madsen 2024). Multitask incentives, agency missions, sequential reporting, and bureaucratic data collection provide related institutional mechanisms (Holmstrom and Milgrom 1991; Dewatripont et al. 1999; Li 2007; Ekmekci et al. 2022; Oechslin and Steiner 2022). The cumulative-information side connects to persistent excitation and information singularity (Ziemann and Sandberg 2021), while sequential experiment choice and sensor scheduling study information acquisition with switching or state costs (Chernoff 1959; Kiefer and Sacks 1963; Whittle 1965; Kolonko and Benzing 1985; Hardwick and Stout 2007; Baras and Bensoussan 1989; Lee et al. 2001; Leong et al. 2016).

Target-weighted posterior loss is the Bayesian A-optimality criterion used in the planner problem (Chaloner and Verdinelli 1995; Pukelsheim 2006). Sequential Bayesian design and active sensing provide related belief-state formulations (Huan et al. 2024; Ahmad et al. 2014), and dynamic information acquisition appears in Zhong (2022); Liang et al. (2022). The empirical acquisition rule is related to costly feature acquisition (Saar-Tsechansky et al. 2009). Policy feedback and measurement-induced behavior provide broader motivation for treating the reporting system as part of the policy environment (Lucas 1976; McLeay and Tenreyro 2020; Chemla and Hennessy 2021). Majorization, matrix analysis, identification, and projection inference supply the technical tools (Marshall et al. 2011; Horn and Johnson 2012; Escanciano 2022; Kaji 2021; Andrews and Mikusheva 2016; Andrews and Guggenberger 2019).

Jehiel and Steiner (2020) study repeated sampling when a decision maker selectively forgets earlier realizations and characterize optimal choice through a second-thought-free condition. The archive here accumulates reports across dates and assigns their future value to a separate policy user. This institutional separation produces the planner–institution wedge and the cumulative-coverage condition.

The paper begins with the reporting institution and then develops the observation experiment, institutional adjustment, archive accumulation, and the planner's reporting choice. The empirical analysis measures the predictive value of retained histories and intermediate reports. Core derivations accompany the paper; dynamic extensions, posterior calculations, data construction, and predictive diagnostics appear in the supplemental appendix.

2. The Reporting Institution

2.1. The Reporting Institution and the Targets

U.S. Call Reports record mortgage composition under different reporting schedules. In March 2017, eligible domestic institutions could adopt form FFIEC 051, which generally used an asset threshold of one billion dollars at introduction (Office of the Comptroller of the Currency, Federal Reserve Board, and Federal Deposit Insurance Corporation 2017). Its adjustable-rate first-lien mortgage item, RCON5370, is reported in June and December.[2] The application changes access to the dated ARM item while retaining the other observed records. The reporting form and item schedule describe the observed installation; retained ARM reports define the archive; subsequent mortgage composition and net interest income define two forecast targets.

The source contains 315,615 bank-quarter observations over 2012–2025, downloaded from the FDIC BankFind API on August 28, 2026 (Federal Deposit Insurance Corporation 2026). Reporting metadata, historical identifiers, mergers, and quarterly financial variables accompany the records. Federal funds and two- and ten-year Treasury rates come from FRED (Federal Reserve Bank of St. Louis 2026). The sample consists of domestic commercial banks on forms 041 or 051 with positive assets and first-lien mortgages, complete common controls, and observed outcomes. Histories restart at mergers, identifier changes, and gaps. The accounting and continuing-bank conditions in Appendix F yield 165,049 evaluated bank-origin observations for 2016–2025. Calculations use the downloaded vintage.[3]

Write A_{it} for ARM balances, T_{it} for first-lien mortgage totals, and S_{it} for bank assets. Put $h_t = 1$ for March and September origins and $h_t = 2$ for June and December origins. The outcomes are

$$Y_{it}^{A} = 100\frac{A_{i,t+h_t}}{T_{i,t+h_t}}, \qquad Y_{it}^{N} = 10^4\frac{\mathrm{INTINQ}_{i,t+1} - \mathrm{EINTXQ}_{i,t+1}}{S_{it}}. \tag{1}$$

Composition is measured in percentage points and income in basis points of origin assets.[4] March and September form-051 ARM entries follow the reporting mask, including positive

[2] The mortgage total and six maturity or repricing buckets remain available quarterly (Federal Financial Institutions Examination Council 2016).

[3] An origin means that quarter's reports are available. The downloaded records omit historical release times and first-release vintages.

[4] Income inputs are quarterly flows, and composition targets fall on June or December reports.

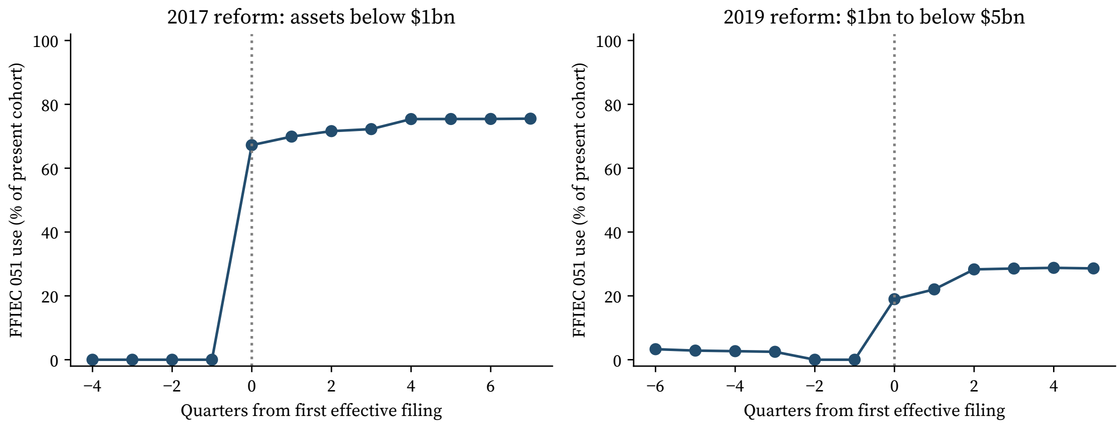

FIGURE 1. Form use after eligibility changes

Cohorts use June 2016 assets for the 2017 reform and June 2018 assets for the 2019 reform. The first effective filings are March 2017 and September 2019. Presence requires the domestic commercial-bank and form conditions described in Appendix J. The vertical line marks the first effective filing. The plotted denominator follows banks present in the cohort sample at each date.

entries found in the downloaded file. Each access comparison holds the bank, origin, other records, and realized outcome fixed.

2.2. Reform Timing and Observed Form Choices

The 2017 introduction removed about 40 percent of the items in form 041 and changed reporting frequencies for selected items (Office of the Comptroller of the Currency, Federal Reserve Board, and Federal Deposit Insurance Corporation 2017). Further revisions effective in June 2018 affected about seven percent of the streamlined form's items (Federal Financial Institutions Examination Council 2018).[5] The ARM item retained its June/December schedule for continuing 051 filers. Form adoption by newly eligible banks is therefore the channel connecting this expansion to their ARM archives.

Figure 1 follows fixed cohorts defined in the preceding June. Among 4,654 domestic commercial banks below one billion dollars in June 2016, 3,015 use form 051 at the first effective filing and 344 first use it later through December 2018. Among 455 banks with June 2018 assets from one billion to below five billion dollars, 81 use 051 in September 2019 and 40 first use it later through December 2020. The plotted denominator is the cohort present within the domestic-commercial-bank form sample at each date. These paths document implementation and choice following the reporting rule. Table 1 follows adjacent observed form choices within the same fixed cohorts. Once a bank reports on 051, the next observed

[5]In September 2019, eligibility expanded to certain domestic institutions below five billion dollars in assets, and approximately 37 percent of first- and third-quarter items were removed from those quarters (Board of Governors of the Federal Reserve System, Federal Deposit Insurance Corporation, and Office of the Comptroller of the Currency 2019).

TABLE 1. Adoption timing and persistence of the installed reporting form

Eligible cohort	Banks	Adopt	Immediate	Retain 051	Enter 051
2017: below $1bn	4,654	74.89	89.76	99.89	4.18
2019: $1–5bn	455	28.34	66.94	99.62	2.59

Percentages use fixed pre-reform asset cohorts. Any adoption divides banks that use FFIEC 051 by banks observed after the effective date. Immediate adoption is the share of adopters first observed on FFIEC 051 at the effective filing. Retention and entry use adjacent observed quarters through 2018Q4 for the 2017 cohort and 2020Q4 for the 2019 cohort. The transition rates describe installed-form histories; bank choice and eligibility conditions determine those histories.

TABLE 2. FFIEC reporting-burden accounting

Filer group	Before	After	Hours saved	Annual dollars	Payback quarters
Existing 051 filers	40.11	39.08	1.03	206	162.1
Newly eligible filers	63.69	50.96	12.73	2,546	13.1

Before, after, and hours saved are average preparation hours per quarterly filing reported in the 2019 final rule. Annual dollars value preparation time at $50 per hour. Payback combines a commenter's estimate of 160 conversion hours and $350 of software cost with the estimated quarterly time saving. The agencies cited variation in training, systems, automation, and vendor costs and left conversion costs unquantified at the agency level. The entries provide an accounting scale for the operating and adjustment costs in the model.

quarter retains that form in 99.89 percent of 2017-cohort transitions and 99.62 percent of 2019-cohort transitions. Entry from 041 occurs in 4.18 and 2.59 percent of adjacent transitions. The observed installation therefore changes through an adoption margin followed by persistence in the selected form. Adoption timing reflects reporting savings, system changes, expected growth, and supervisory requirements. The 2019 rule estimates average quarterly preparation time of 40.11 to 39.08 hours for existing 051 filers and 63.69 to 50.96 hours for newly eligible filers (Board of Governors of the Federal Reserve System, Federal Deposit Insurance Corporation, and Office of the Comptroller of the Currency 2019).[6] Table 2 combines these figures at a stated labor price. The resulting payback differs by more than a factor of twelve across filer groups. These whole-form figures anchor the flow and adjustment costs associated with reporting changes. Item-level ARM costs and shared preparation costs enter the acquisition decision separately.

The FFIEC began its 2027 full review of Call Reports in October 2025 and collected industry comments on reporting burden and data use through January 2026. Its 2025 annual report places any resulting Paperwork Reduction Act notices by the end of 2026 (Federal Financial Institutions Examination Council 2026). This review preserves the

[6]The 2017 notice also records a commenter's estimate of 160 conversion hours, primarily for training, and $350 for software. The agencies describe heterogeneity in systems, automation, training, and vendor charges (Office of the Comptroller of the Currency, Federal Reserve Board, and Federal Deposit Insurance Corporation 2017).

policy setting in which reporting content, filer resources, and subsequent data use are considered together.

The local first stage is small in 2017 and zero in the 2019 window. Within ten percent of the 2017 threshold, nine of 59 present below-threshold banks use 051 in March, compared with zero of 58 above. Within ten percent of the 2019 threshold, 16 below and ten above are present in September; both sides have zero adopters, also in March 2020. Appendix J records the eligibility and contemporaneous-rule qualifications. The application uses these histories as institutional evidence about reporting schedules and estimates access values on fixed realized outcomes.

3. Institutions and the observation experiment

The observation experiment describes what a report can reveal about the state. The operating objective determines which direction the institution chooses to measure. Later archive users value that information through their own targets. The following notation keeps the information technology fixed while the direction of the report changes; the two-direction case then gives the policy mechanism used throughout the main analysis.

3.1. The reporting technology

Let $\theta \in \mathbb{R}^2$ contain two economic coefficients in fixed units. A report measures a linear combination of them. The organization chooses a response direction α_t, the reporting agency chooses a measurement direction β_t, and their difference determines which combination the report reveals. Signal capacity and measurement noise are held fixed as these installed directions change. A general response-and-measurement system embeds this two-direction experiment and tracks changes of units.

3.2. Operating priorities

Operating values and exposures determine desired reporting directions. The dynamic model takes the resulting angular priority path as given and assigns costs to departures from it and to changes in installed direction. The downstream user's information loss enters the planner's objective with weight χ; this allocation of objectives specifies the externality. The report reaches an authorized archive user under the maintained access arrangement.[7] The archive problem here values a retained record for the stated future decision under fixed access and operating costs.

[7] Public dissemination is a further institutional choice that can change current incentives, trading, or the information acquired by other agents.

3.3. Two Directions

For the two economic coordinates, write

(2)

$$Q(x) = \begin{pmatrix} \cos x & -\sin x \\ \sin x & \cos x \end{pmatrix}, \qquad e(x) = \begin{pmatrix} \cos x \\ \sin x \end{pmatrix}, \qquad V_t = Q(\alpha_t), \quad U_t = Q(\beta_t), \quad \delta_t = \beta_t - \alpha_t.$$

Let $g > 0$ be the fixed response scale and $\Lambda = \mathrm{diag}(\lambda_H, \lambda_L)$ the fixed measurement capacities, with $\lambda_H > \lambda_L \geq 0$. One observation supplies

$$\mathcal{I}(\delta) = g^2 Q(\delta)\Lambda Q(\delta)' = g^2\{\lambda_L I_2 + (\lambda_H - \lambda_L)e(\delta)e(\delta)'\}. \tag{3}$$

An exact experiment yielding this matrix is $Y = g\Lambda^{1/2}Q(\delta)'\theta + \varepsilon$, with $\varepsilon \sim N(0, I_2)$. The second component carries zero signal when $\lambda_L = 0$. The signs of the bases are fixed by an installation convention. Angles in the adjustment problem are real-valued installed coordinates; observation precision is periodic with period π.

The information matrix in (3) describes a fresh observation. Retained records contribute to the archive used for the policy decision. Target costs, economic-unit transformations, and the Schur–Horn allocation set describe feasible fresh observations. Reporting paths and policy choices use the accumulated posterior precision.

4. Intertemporal Institutional Adjustment

4.1. Objectives and Timing

At date t, an organization inherits α_t and chooses α_{t+1}. An agency inherits β_t and chooses β_{t+1}. The next report uses the new directions. Their desired positions are $\bar{\alpha}_t$ and $\bar{\beta}_t$, with $\bar{\delta}_t = \bar{\beta}_t - \bar{\alpha}_t$. For either decision maker, write the installed coordinate as x_t and its desired coordinate as $\bar{x}_t$. The problem is

$$\min_{\{x_{t+1}\}_{t\geq 0}} \sum_{t=0}^{\infty} \eta^t \left\{\frac{\rho}{2}(x_{t+1} - \bar{x}_t)^2 + \frac{\kappa}{2}(x_{t+1} - x_t)^2\right\}, \qquad \rho, \kappa > 0, \quad 0 < \eta < 1. \tag{4}$$

The priority sequence is bounded and known at date zero. A constant sequence describes a regime expected to remain in force. Information about later changes can be incorporated by resolving the same problem after an announcement. Operating costs measure resources or losses that also enter the planner's objective. The parameters ρ, κ, η are common across the two institutions in the results using δ_t as their installed state.

4.2. Adjustment with Continuation Value

LEMMA 1 (Quadratic institutional adjustment). *Let $P > 0$ solve*

$$P = \frac{\kappa(\rho + \eta P)}{\rho + \kappa + \eta P}, \qquad \gamma = \frac{\kappa}{\rho + \kappa + \eta P}. \tag{5}$$

For $0 < \eta < 1$, $P = \{\sqrt{[\rho + \kappa(1-\eta)]^2 + 4\eta\kappa\rho} - [\rho + \kappa(1-\eta)]\}/(2\eta)$. *The unique optimal path satisfies*

$$x_{t+1} = \gamma x_t + (1-\gamma)\widetilde{x}_t, \qquad \widetilde{x}_t = (1-\eta\gamma)\sum_{j=0}^{\infty}(\eta\gamma)^j \bar{x}_{t+j}. \tag{6}$$

Consequently $\delta_{t+1} = \gamma\delta_t + (1-\gamma)\widetilde{\delta}_t$, *where* $\widetilde{\delta}_t$ *is the same weighted average of* $\bar{\delta}_{t+j}$. *For a constant priority* $\bar{\delta}$, $\delta_t - \bar{\delta} = \gamma^t(\delta_0 - \bar{\delta})$. *Holding* ρ *and* η *fixed,* γ *increases with* κ.

The Bellman equation yields the minimizing path. The continuation term ηP enters the denominator in (5). The case $\eta = 0$ gives the one-period coefficient $\kappa/(\rho + \kappa)$. For $\eta > 0$, anticipated priorities enter today's installation through the weights in (6).

An announced unit change confined to priority date $t + j$ changes x_{t+1} by $(1-\gamma)(1-\eta\gamma)(\eta\gamma)^j$, holding x_t fixed. A constant priority following a new announcement produces a geometric transition from the inherited position. These responses separate announcement timing, installation, and the date at which the resulting report enters the archive.

When operating or adjustment curvatures differ across institutions, each direction follows its own equation (6). The installed state is then (α_t, β_t). With the common coefficients assumed above, the relative direction itself obeys the displayed scalar transition.

5. Data Histories and Policy Learning

5.1. The Archive and the Policy Decision

Assume a fixed parameter $\theta \sim N(\mu_0, J_0^{-1})$, where $J_0 \succ 0$. At reporting date s, n_s conditionally independent observations follow the experiment in (3). Their noises are independent across dates. The records include the direction, sample size, and measurement map for every date, and the policymaker retains all observations. The posterior precision after reports $0, \ldots, T-1$ is

$$J_T = J_0 + \sum_{s=0}^{T-1} n_s \mathcal{I}(\delta_s). \tag{7}$$

The report at date zero can be inherited. The institutions subsequently choose the directions of reports $1, 2, \ldots$ by (6). Reindexing the archive by one period leaves the results unchanged.

After observing a policy problem (a, c), with $c > 0$, the policymaker chooses $q \in \mathbb{R}$ to minimize $L(q, \theta; a, c) = cq^2/2 - qa'\theta$. Given θ, the minimizing action is $q^*(\theta) = a'\theta/c$. The distribution of (a, c) is independent of the parameter and observations, and $D = E[aa'/c]$ is finite. It may depend on a known policy regime.

LEMMA 2 (Bayesian quadratic decision loss). *Let μ_T be the posterior mean. The Bayes action is $\widehat{q} = a'\mu_T/c$, and its expected loss relative to $q^*(\theta)$ is*

$$\mathcal{R}(J_T; D) = E[L(\widehat{q}, \theta; a, c) - L(q^*(\theta), \theta; a, c) \mid J_T] = \frac{1}{2}\operatorname{tr}(DJ_T^{-1}). \tag{8}$$

For fixed $\mathcal{I} \succ 0$, J_0, and D, $\mathcal{R}(J_0 + n\mathcal{I}; D) = \operatorname{tr}(D\mathcal{I}^{-1})/(2n) + O(n^{-2})$.

The equality uses a proper prior and the Gaussian experiment. In the expansion, $\mathcal{I}$ is fixed and positive definite. The target cost in (37) describes the coefficient of $1/n$ under those conditions. The exact loss (8) applies throughout the Gaussian specification, including rank-one reports.

5.2. Adjustment Costs and the Information Left by Specialization

Suppose the priority remains at $\bar{\delta}$. Rotate coordinates so that $u = e(\bar{\delta})$ is the first coordinate and $v = e(\bar{\delta} + \pi/2)$ is the second. The rank-one case, $\lambda_L = 0$, makes each report informative about its measured linear combination. A finite set of nonparallel reports identifies both coefficients in the associated linear mean experiment. The rate at which new directional information accumulates determines posterior concentration.

LEMMA 3 (Cumulative excitation). *In the fixed-state Gaussian experiment with deterministic reporting maps and $J_0 \succ 0$, the posterior covariance converges to zero if and only if $\lambda_{\min}(\sum_{s<T} n_s \mathcal{I}(\delta_s)) \to \infty$. A uniform lower bound $\sum_{s=t}^{t+L-1} n_s \mathcal{I}(\delta_s) \succeq \epsilon I_d$ for some fixed $L \geq 1$, $\epsilon > 0$, and every t implies covariance of order $O(T^{-1})$.*

The first condition concerns cumulative information in every direction. The second is a finite-window persistent-excitation condition of the type studied in adaptive control (Boyd and Sastry 1983; Bai and Sastry 1985).[8] Sparse supplemental reports can satisfy the cumulative condition with increasingly long gaps between them. Eigenvalue bounds for J_T give the result.

[8]The finite-window condition supplies a uniform lower rate of information accumulation. Divergent cumulative information alone permits long intervals during which a neglected direction receives no new precision.

THEOREM 1 (Specialization and incomplete learning). *Let* $n_s = n > 0$*,* $\lambda_L = 0$*,* $\tau = ng^2\lambda_H$*, and* $\delta_s = \bar{\delta} + \gamma^s \varepsilon_0$ *with* $0 < \gamma < 1$*. Define*

$$C_\infty = v' J_0 v + \tau \sum_{s=0}^{\infty} \sin^2(\gamma^s \varepsilon_0). \tag{9}$$

Then

$$J_T^{-1} \longrightarrow \frac{vv'}{C_\infty}, \qquad \mathcal{R}(J_T; D) \longrightarrow \frac{v'Dv}{2C_\infty}, \qquad C_\infty \le v'J_0 v + \frac{\tau\varepsilon_0^2}{1-\gamma^2}. \tag{10}$$

For $0 < |\varepsilon_0| \le \pi/2$*,* C_∞ *is strictly increasing in* γ*. For* $v'Dv > 0$*, the limiting policy loss therefore decreases with* κ*, holding the initial gap, priorities,* ρ*,* η*, and experiment fixed.*

The path generated by adjustment contributes a finite amount of precision about the coefficient orthogonal to the eventual reporting direction. Increasing the adjustment cost lengthens this transition and preserves additional observations of that coefficient. This statement compares information along the decentralized paths. A welfare comparison also includes the operating and adjustment costs in (4).

The proof writes J_T in the (u, v) basis. Its first diagonal entry diverges linearly in T, its off-diagonal entry converges, and its second diagonal entry converges to C_∞. Matrix inversion gives (10). The convergence of the off-diagonal entry uses the geometric rate as well as the boundedness of the angular gap.

To compare specialization speeds within the same technology, consider geometric paths $\delta_t - \bar{\delta} = \gamma^t \varepsilon_0$ for $\gamma \in [0, 1]$. The discounted resource cost associated with the relative direction is

$$\mathcal{C}^G(\gamma) = \frac{\varepsilon_0^2}{4} \frac{\rho\gamma^2 + \kappa(1-\gamma)^2}{1-\eta\gamma^2}. \tag{11}$$

For $\bar{d} = v'Dv > 0$, define $\mathcal{L}^G(\gamma) = \bar{d}/[2C_\infty(\gamma)]$ on $[0, 1)$ and $\mathcal{L}^G(1) = 0$. At $\gamma = 1$, the installed direction remains at its initial value and orthogonal precision diverges.

PROPOSITION 1 (The planner's specialization speed). *Let* $0 < |\varepsilon_0| \le \pi/2$*,* $\rho, \kappa > 0$*,* $\eta \in (0, 1)$*, and* $\bar{d} > 0$*. The decentralized speed* γ^E *uniquely minimizes* $\mathcal{C}^G$ *and satisfies*

$$\rho\gamma^E = \kappa(1-\gamma^E)(1-\eta\gamma^E). \tag{12}$$

For every $\chi > 0$*, each global minimizer of* $\mathcal{C}^G(\gamma) + \chi\mathcal{L}^G(\gamma)$ *lies in* $(\gamma^E, 1]$*. For all sufficiently small* $\chi > 0$*, this minimizer is unique and interior, and*

$$\gamma^{SP} = \gamma^E + \chi \frac{\bar{d} C_\infty'(\gamma^E)}{2C_\infty(\gamma^E)^2 \mathcal{C}^{G\prime\prime}(\gamma^E)} + O(\chi^2) > \gamma^E. \tag{13}$$

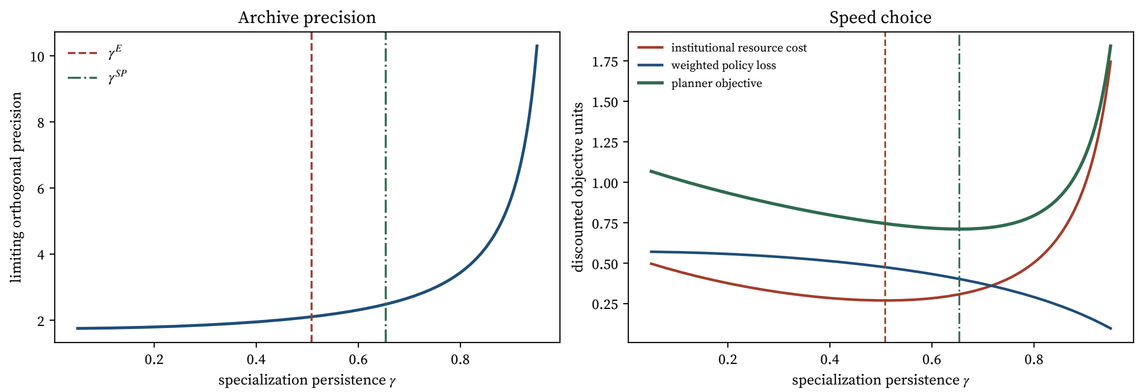

FIGURE 2. Specialization speed and the archive externality

The calculation uses $\rho = 1$, $\kappa = 2$, $\eta = 0.95$, $\varepsilon_0 = \pi/3$, $v' J_0 v = \tau = \bar{d} = 1$, and $\chi = 2$. The decentralized choice is $\gamma^E = 0.508$ and the planner's choice is $\gamma^{SP} = 0.654$. The left panel evaluates (9); the right panel evaluates (11), weighted limiting policy loss, and their sum.

The operating institution chooses γ^E from current operating and reconfiguration costs. The planner values the precision accumulated during the transition and delays specialization. Equation (13) measures the local speed wedge in policy demand, archive sensitivity, and the curvature of resource cost.

COROLLARY 1 (Complete-learning characterization with sparse reports). *If* $\lambda_L > 0$ *and* $n_s = n$, *then* $\mathcal{R}(J_T; D) \leq \operatorname{tr}(D)/(2Tng^2\lambda_L)$ *for* $T \geq 1$. *Under the rank-one path of Theorem 1, let* m_T *be the cumulative number of independent supplemental reports of* $v'\theta$ *available by date* T, *each with precision* $\xi > 0$. *The posterior covariance tends to zero if and only if* $m_T \to \infty$. *Complete learning therefore permits a vanishing reporting frequency, including any diverging sequence with* $m_T/T \to 0$.

The two statements characterize two routes to posterior concentration. A positive per-report precision floor supplies information in every direction at every date. Under specialization, complete learning requires and is achieved by unbounded cumulative precision in the direction left by the convergent installation path; its arrival rate can vanish.

5.3. Changes in the Parameter or the Archive

Equation (7) uses a stable coefficient vector and retained records with known observation maps. If the coefficient follows $\theta_{t+1} = \theta_t + \omega_{t+1}$, with independent Gaussian innovation covariance $Q_t \geq 0$, prediction and updating give

$$J_{t+1}^{-} = (J_t^{-1} + Q_t)^{-1}, \qquad J_{t+1} = J_{t+1}^{-} + n_{t+1}\mathcal{I}(\delta_{t+1}). \tag{14}$$

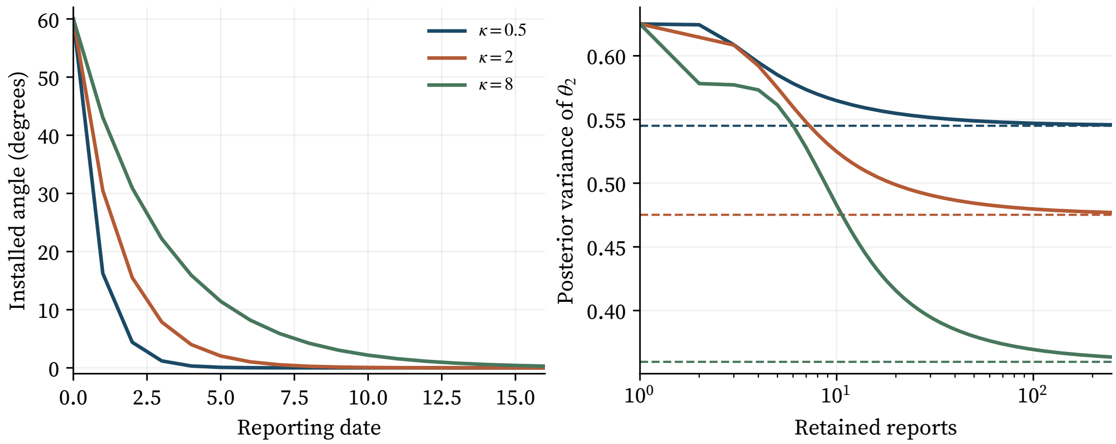

FIGURE 3. Installation speed and posterior uncertainty

Computed examples use $\rho = 1$, $\eta = 0.95$, $\kappa \in \{0.5, 2, 8\}$, $J_0 = I_2$, $\tau = 1$, $\bar{\delta} = 0$, and $\varepsilon_0 = \pi/3$. The left panel plots the optimal installed direction. The right panel plots posterior variance of the second coefficient; dashed lines are the limits in Theorem 1. Values are generated by the Gaussian model.

The decision loss remains (8) for the current posterior. Theorem 1 uses $Q_t = 0$. A retained subset of records contributes its own sum of information matrices. Loss of regime labels changes the likelihood itself;[9] Labelled and unlabelled observations therefore carry different information. These specifications make parameter drift, record retention, and map retention explicit inputs to the learning problem.

5.4. An Additional Record and a Future Target

The archive can also be represented by its posterior covariance P_t. For a Gaussian state with transition $\theta_{t+1} = A\theta_t + \omega_{t+1}$ and independent innovation covariance Q, prediction gives $P_{t+1}^{-} = AP_tA' + Q$. A report $r = s'\theta_t + \varepsilon$, with independent variance $\nu \geq 0$, updates this covariance to $P_t - P_t ss' P_t/(\nu + s'P_t s)$ whenever the denominator is positive. These are the covariance recursions of Kalman (1960). They allow a current coordinate to be observed exactly while leaving uncertainty about another coordinate, such as its subsequent rate of change.

LEMMA 4 (Archive-conditional value of a record). *Suppose the report and the subsequent state innovations satisfy the Gaussian specification above, and acquiring the report leaves the state transition unchanged. For a policy decision h dates later with loss matrix $D \geq 0$, the reduction in expected posterior decision loss is*

$$\Delta_h(P_t; s, \nu, D) = \frac{s'P_t(A^h)'DA^hP_ts}{2(\nu + s'P_ts)}. \tag{15}$$

[9] Discarding a record removes one likelihood contribution. Discarding its regime label replaces a labelled likelihood by a mixture, so its information loss generally differs from deleting the same matrix term.

The value is zero exactly when $D^{1/2}A^h P_t s = 0$. For a scalar forecast target $u'\theta_{t+h}$ under squared error, its expected MSE reduction is (15) *with $D = uu'$ and the factor* 1/2 *removed.*

The numerator measures how uncertainty in the report covaries with uncertainty in the future target after propagation through A^h. The denominator measures the report's predictive variance. Thus two histories that imply different P_t can give different values to the same report at the same current installation. A target whose residual uncertainty has little covariance with the report has little expected gain. The application constructs such an uncertainty state, evaluates its forecast calibration, and compares the predicted gain with losses observed after an additional record becomes available.

5.5. State Persistence and Effective Archive Age

The fixed-state archive in Theorem 1 retains every precision contribution. A changing state discounts earlier records through the transition matrix in Lemma 4. The following result separates the value of record age from the steady filtering loss created by a specialized reporting system.

PROPOSITION 2 (State persistence and a specialized archive). *Suppose $\theta_{t+1} = \phi\theta_t + \omega_{t+1}$, where $|\phi| < 1$ and $\omega_{t+1} \sim N(0, qI_2)$ with $q > 0$. First, for any report described in Lemma 4,*

$$\Delta_h(P_t; s, v, D) = \phi^{2h}\Delta_0(P_t; s, v, D), \qquad \sum_{h=0}^{H} \eta^h \Delta_h = \Delta_0 \frac{1 - (\eta\phi^2)^{H+1}}{1 - \eta\phi^2}. \tag{16}$$

Suppose that from date $T + 1$ onward every report has rank-one precision $\tau e_1 e_1'$, with $\tau > 0$, and that $P_T = \operatorname{diag}(p_{1,T}, p_{2,T})$. The posterior covariance remains diagonal and converges to $P^* = \operatorname{diag}(p_1^*, p_2^*)$, where

$$p_2^* = \frac{q}{1-\phi^2}, \qquad p_1^* = \begin{cases} \dfrac{\sqrt{(1+\tau q-\phi^2)^2 + 4\tau\phi^2 q} - (1+\tau q-\phi^2)}{2\tau\phi^2}, & \phi \neq 0, \\[2ex] \dfrac{q}{1+\tau q}, & \phi = 0. \end{cases} \tag{17}$$

For two diagonal covariances P_T and $\widetilde{P}_T$ followed by the same specialized reports,

$$|p_{1,T+h} - \widetilde{p}_{1,T+h}| \le |\phi|^{2h}|p_{1,T} - \widetilde{p}_{1,T}|, \qquad |p_{2,T+h} - \widetilde{p}_{2,T+h}| = |\phi|^{2h}|p_{2,T} - \widetilde{p}_{2,T}|. \tag{18}$$

Consequently, for $D \succeq 0$, the posterior policy-loss difference obeys

$$|\mathcal{R}(P_{T+h}; D) - \mathcal{R}(\widetilde{P}_{T+h}; D)| \le \frac{|\phi|^{2h}}{2}\{D_{11}|p_{1,T} - \widetilde{p}_{1,T}| + D_{22}|p_{2,T} - \widetilde{p}_{2,T}|\}. \tag{19}$$

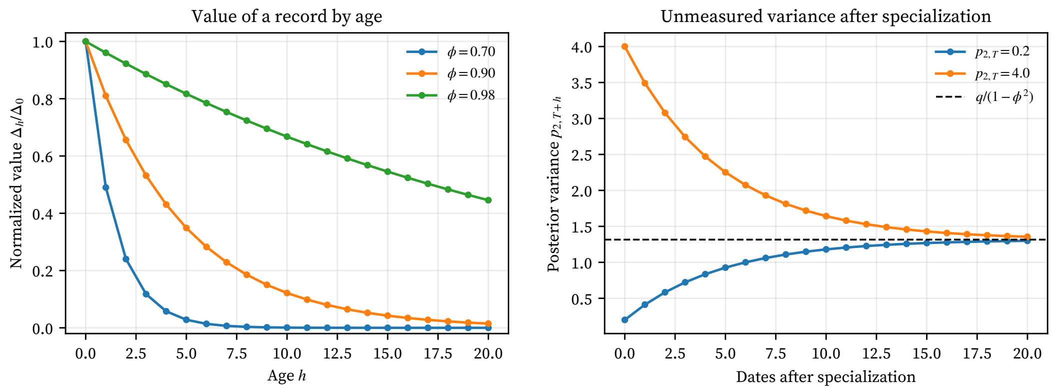

FIGURE 4. State persistence and the effective age of a record

The left panel plots the normalized value $\Delta_h/\Delta_0 = \phi^{2h}$. The right panel plots the unmeasured-coordinate variance after specialization for $q = 0.25$, $\phi = 0.9$, and two inherited covariances. Its limit is $q/(1-\phi^2)$. Values follow Proposition 2.

Equation (16) gives a record an economic half-life determined by state persistence. Equation (17) gives the current-state uncertainty produced by permanent specialization: the measured coordinate reaches a scalar filtering variance, while the other coordinate reaches its unconditional stationary variance. Reports collected during installation change finite-horizon losses according to (18). Their effect lasts longer when $|\phi|$ is closer to one.

Theorems 1 and 2 describe two archive regimes. A stable coefficient preserves the directional precision collected during institutional transition. A stationary changing state depreciates that precision and makes recent directional coverage enter current policy loss. Adjustment cost then affects the timing of information benefits, while state persistence determines how much of each benefit reaches later users. This distinction matches the application, where removing recent reports changes composition forecasts more than adding older records.

6. The Planner's Reporting Choice

6.1. A State Consisting of Installation and Precision

The planner includes the costs of both institutions and assigns weight $\chi \geq 0$ to policy losses. Write $m = (\alpha + \beta)/2$. Their combined one-period cost separates into a component for the common direction and

$$C_t(x, y) = \frac{\rho}{4}(y - \bar{\delta}_t)^2 + \frac{\kappa}{4}(y - x)^2, \qquad x = \delta_t, \quad y = \delta_{t+1}. \tag{20}$$

The common direction follows its own quadratic problem. Information depends on y, so its contribution can be removed from the planner's relative-direction problem. For dates

$t = 0, \ldots, H-1$, let the known sequence n_t describe the next reporting batches and let D_t describe the policy problems following those batches. With $\mathcal{V}_H = 0$, the planner's value is

$$
\begin{aligned}
\mathcal{V}_t(x, J) = \min_{y \in \mathbb{R}} \Big\{ & C_t(x, y) + \frac{\chi}{2} \operatorname{tr}\Big(D_t [J + n_t \mathfrak{I}(y)]^{-1}\Big) \\
& + \eta \mathcal{V}_{t+1}\big(y, J + n_t \mathfrak{I}(y)\big) \Big\}.
\end{aligned}
\tag{21}
$$

In (21), J is the precision available before the next reporting batch. With fixed target demand and Gaussian posterior covariance invariant to report realizations, the planner's state is (x, J) together with the date and exogenous schedule. The posterior mean determines the policy action and integrates out of expected regret. The objective allocates institutional costs and the benefits received by later policy users, with χ expressing that allocation in the units of operating cost.

Conditional on the initial archive, a sequence of chosen angles determines the entire precision path in the baseline Gaussian specification. A date-zero reporting schedule then implements the planner allocation, while realized report values determine subsequent policy actions through the posterior mean. With a latent target regime, the first report updates both the Gaussian coefficient posterior and the regime posterior. The resulting recursion has state (x, μ, P, π), and the realized report changes (μ, π) while its precision changes P. Two states with the same installation and coefficient precision can therefore induce different next reports.

For finite H, positive J, and finite D_t, an optimal choice exists. The quadratic adjustment term makes the objective coercive in y. Moreover, $\mathcal{V}_t(x, J)$ decreases in the Loewner order on J. At a differentiable branch of the value function, define $J^+ = J + n_t \mathfrak{I}(y)$ and

$$
M_t = \frac{\chi}{2} (J^+)^{-1} D_t (J^+)^{-1} - \eta \nabla_J \mathcal{V}_{t+1}(y, J^+). \tag{22}
$$

The matrix M_t is positive semidefinite. An interior optimum satisfies

$$
\frac{\rho}{2}(y - \bar{\delta}_t) + \frac{\kappa}{2}(y - x) + \eta \partial_x \mathcal{V}_{t+1}(y, J^+) = n_t \operatorname{tr}\{M_t \mathfrak{I}'(y)\}. \tag{23}
$$

The right side is the marginal value of changing the next batch's information. It includes the effects on the current policy decision and on subsequent uses of the archive. Its derivative requires the stated differentiability qualification.

Report contents and subsequent targets. A content-dependent extension allows a report to change beliefs about both coefficients and the relevance of a later target. This adds a threshold indexed by the report realization.

6.2. Institutional Costs and Benefits to Archive Users

For the finite-horizon institutional problem, let C_H be the discounted sum of the two institutions' resource costs after removing the common direction, and let R_H be the discounted sum of posterior policy losses. The decentralized path δ^E minimizes C_H; the planner's path δ^{SP} minimizes $C_H + \chi R_H$. Their objective difference is

$$\mathcal{W}_H = \chi\{R_H(\delta^E) - R_H(\delta^{SP})\} - \{C_H(\delta^{SP}) - C_H(\delta^E)\} \geq 0. \tag{24}$$

Current institutions bear the second term, and subsequent users receive the first. The inequality follows by including δ^E in the planner's feasible set. It measures the welfare wedge in the units specified by χ and the resource costs. Under the constant-priority assumptions of Theorem A2, $\delta^E = 0$ and $\mathcal{W}_H > 0$ precisely above its threshold. A comparison across values of κ changes the underlying resource technology as well as reporting behavior; equation (24) compares paths at a fixed technology.

PROPOSITION 3 (Implementation by archive-value payment). *Suppose an archive user can commit before reporting to a path-contingent payment*

$$T_H(\delta) = \chi\{R_H(\delta^E) - R_H(\delta)\}. \tag{25}$$

An institution choosing a feasible path to minimize $C_H(\delta) - T_H(\delta)$ selects the planner's path. The payment is zero at δ^E. At every planner solution distinct from δ^E, it is positive and the archive user's gross decision-loss reduction equals the payment. The net social gain equals (24) *because the payment is a transfer.*

PROOF. The institution's net cost is $C_H(\delta) + \chi R_H(\delta) - \chi R_H(\delta^E)$, so its minimizers equal the planner's minimizers. Strict convexity of C_H gives $C_H(\delta^{SP}) > C_H(\delta^E)$ whenever $\delta^{SP} \neq \delta^E$. Planner optimality then implies

$$T_H(\delta^{SP}) \geq C_H(\delta^{SP}) - C_H(\delta^E) > 0. \tag{26}$$

Subtracting the resource-cost increase from this payment gives (24). □

The payment uses quantities already present in the planner's problem: declared policy targets, inherited precision, reporting maps, and institutional resource costs. A reporting standard can implement the same allocation by restricting the feasible path to a planner solution. In the payment formulation, the information externality is the reduction in policy loss created for later users.

6.3. Inherited Precision and the Next Report

The terminal reporting choice yields a solved case of (21). Set $x = \bar{\delta} = 0$, $h = \rho + \kappa$, and let inherited precision be $J = \mathrm{diag}(j_1, j_2)$, with $j_1 \geq j_2 > 0$. A batch of n reports contributes

$$n\mathcal{I}(\delta) = c_0 I_2 + \tau e(\delta)e(\delta)', \qquad c_0 = ng^2\lambda_L, \quad \tau = ng^2(\lambda_H - \lambda_L) > 0. \tag{27}$$

Use $D = I_2$, which can be generated by policy target $a = \sqrt{2}e_1$ or $\sqrt{2}e_2$ with equal probability and $c = 1$. Define $a_0 = j_1 + c_0$, $b_0 = j_2 + c_0$, and

$$K = b_0(a_0 + \tau), \qquad L = \tau(a_0 - b_0), \qquad N = a_0 + b_0 + \tau. \tag{28}$$

The planner minimizes

$$F(\delta) = \frac{h}{4}\delta^2 + \frac{\chi}{2}\frac{N}{K + L\sin^2\delta}. \tag{29}$$

For any information-equivalent angle, the planner selects a representative with $|\delta| \leq \pi/2$ because it has the smallest installation displacement from zero.

THEOREM 2 (Archive-dependent reporting threshold). *If $a_0 = b_0$, the unique minimizer of (29) is zero. If $a_0 > b_0$, let*

$$\chi^* = \frac{hK^2}{2NL} = \frac{hb_0^2(a_0 + \tau)^2}{2\tau(a_0 - b_0)(a_0 + b_0 + \tau)}. \tag{30}$$

For $0 \leq \chi \leq \chi^*$, the unique global minimizer is zero. For $\chi > \chi^*$, the global minimizers are $\pm\delta^*$, where $\delta^* \in (0, \pi/2)$ is the unique positive solution of

$$h\delta^* = \frac{\chi NL\sin(2\delta^*)}{[K + L\sin^2\delta^*]^2}. \tag{31}$$

The magnitude δ^* increases with χ/h, tends to zero as $\chi \downarrow \chi^*$, and tends to $\pi/2$ as $\chi \to \infty$. Holding b_0, τ, h fixed, χ^* strictly decreases with a_0, from $+\infty$ as $a_0 \downarrow b_0$ to $hb_0^2/(2\tau)$ as $a_0 \to \infty$.

Both signs rotate toward the other coordinate. Their equality reflects the reflection symmetry of the specified archive, policy demand, and operating target. The threshold in (30) is a global statement. Its proof uses the exact objective difference

$$F(\delta) - F(0) = \frac{h}{4}\delta^2 - \frac{\chi NL\sin^2\delta}{2K(K + L\sin^2\delta)}. \tag{32}$$

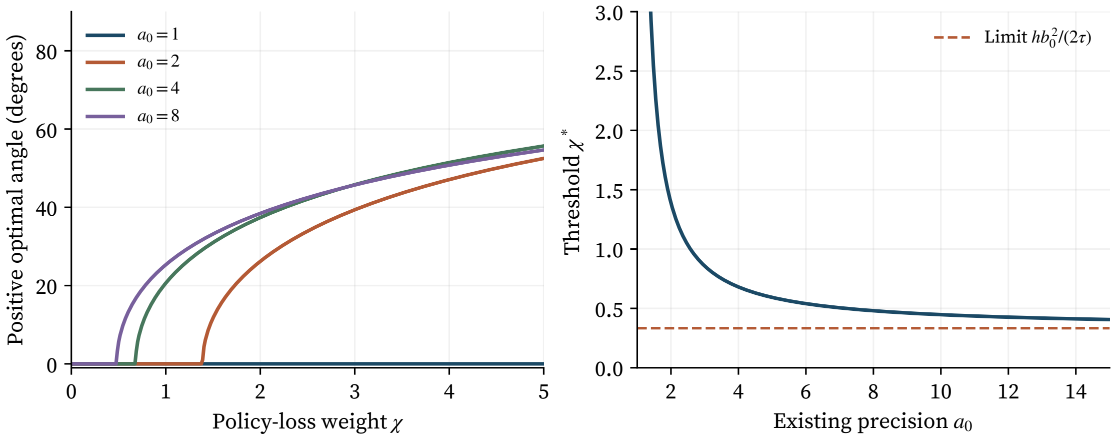

FIGURE 5. The next reporting angle conditional on the archive

Computed examples use $b_0 = 1$, $\tau = 3$, and $h = 2$. The left panel plots the positive optimal angle against χ for $a_0 \in \{1, 2, 4, 8\}$. The right panel plots the threshold in (30); the dashed line is $hb_0^2/(2\tau)$. A negative optimum of equal magnitude exists above each threshold.

The inequality $\sin^2\delta \leq \delta^2$ verifies optimality of zero below the threshold. Above it, the curvature at zero is negative and the positive stationary equation has one solution. The stationary equation also gives the derivative of the threshold.

COROLLARY 2 (Operating continuation after the policy decision). *Suppose the planner chooses one report, policy loss is incurred after that report, and the institutions subsequently solve* (4) *with constant zero priorities. Starting from* $\alpha = \beta = 0$, *the reporting choice is characterized by Theorem 2 with* h *replaced by* $h_c = \rho + \kappa + \eta P$, *where* P *solves* (5).

The continuation term prices the subsequent operating and adjustment costs caused by the intervention. It follows from summing the two institutions' value functions and removing their common direction.

The archive creates the directional policy incentive in this example. When $j_1 = j_2$, isotropic demand values every rotation of the new batch equally. With $j_1 > j_2$, a rotation toward the second coordinate increases the determinant of posterior precision while preserving its trace. The inverse trace in (29) then falls. Additional historical precision about the first coefficient lowers the policy-loss weight at which paying for the rotation becomes worthwhile. The adjustment cost remains $h\delta^2/4$ throughout this comparison.

6.4. A Local Reporting Change and Its Value

For a general terminal state, completing the square in (20) gives $h(\delta - \delta^E)^2/4$ plus a constant, where $\delta^E = (\rho\bar{\delta} + \kappa x)/h$. Define $r(\delta) = \mathrm{tr}\{D[J + n\mathcal{I}(\delta)]^{-1}\}/2$. For a branch of minimizers converging to δ^E as $\chi \to 0$, the implicit function theorem gives

$$\delta^{SP} - \delta^E = -\frac{2\chi}{h} r'(\delta^E) + O(\chi^2), \qquad F(\delta^E) - F(\delta^{SP}) = \frac{\chi^2}{h}[r'(\delta^E)]^2 + O(\chi^3). \tag{33}$$

Here $r'(\delta) = -n\,\mathrm{tr}\{J_\delta^{-1} D J_\delta^{-1} \mathcal{I}'(\delta)\}/2$, with $J_\delta = J + n\mathcal{I}(\delta)$. The local change depends on the precision already available and on policy demand. At the symmetric state in Theorem 2, $r'(0) = 0$. The global comparison in that theorem supplies the positive threshold and the subsequent reporting choices.

Reporting sequences. Multi-date reporting yields a sequence of thresholds and a straddling case. The baseline threshold holds the stated archive, operating costs, and target fixed.

7. Archive value in U.S. bank reports

7.1. The archive and the evaluation target

The theory assigns value to a report through its contribution to a later target. This exercise measures that contribution as predictive loss reduction in U.S. bank records. The pre-report state consists of retained mortgage-share histories and current controls, and the acquisition rule is chosen before the intervening report is revealed. Gaussian and point-mass/change specifications describe that report. Their estimation chronology and calibration are in Online Appendix M. Realized paired forecast losses compare rankings at a common report budget. Item costs and the conversion from predictive loss to welfare remain separate inputs to the planner's objective.

7.2. Valuing a Report Before Its Contents Are Supplied

Start from up to twelve June and December reports and consider acquiring the current March or September report. Every baseline record is retained, so an archive with twelve entries grows to thirteen. The Gaussian-increment model gives the ex ante score

$$g_{it}^G = \widehat{q}_{it}\left\{d_{it} + h_t + \frac{(d_{it} + h_t)^2}{\lambda_{it}} - h_t - \frac{h_t^2}{\lambda_{it} + d_{it}}\right\}. \tag{34}$$

The score uses the semiannual archive and the known horizon before the current report value is supplied. Its drift-uncertainty term accounts for 6.26 percent of total fitted gain, and innovation scale accounts for most of its ranking. The mixed predictive state adds the probability that the intermediate report repeats the latest share, the conditional mean and variance of a change, the implied report variance, and g_{it}^G to the pre-report archive and current controls.

TABLE 3. Report-budget frontier for downstream forecast loss

Budget	Uniform	Innovation	Mixed value	Δ vs. uniform	Δ vs. innovation
10%	4.33%	14.98%	12.98%	3.962 [1.521,6.403]	-0.918 [-2.563,0.726]
25%	10.86%	22.26%	25.05%	6.501 [4.283,8.719]	1.278 [-1.569,4.126]
50%	21.73%	31.44%	33.84%	5.545 [2.719,8.370]	1.097 [-1.474,3.668]

The first three value columns are percentage reductions in the equal-weight mean of ridge, persistence, and gradient-boosting squared forecast loss during 2018–2025. The last two columns report mixed-value minus benchmark gains per eligible origin in squared percentage points. Brackets are simultaneous 95 percent intervals over the three displayed budgets with bank and target-quarter dependence.

For forecast procedure $m \in \{R, P, B\}$, let $\widehat{y}^m_{it,S}$ use the semiannual archive and let $\widehat{y}^m_{it,S+}$ also use the intermediate report. Define its realized paired gain and the acquisition target by

$$G^m_{it} = (Y^A_{it} - \widehat{y}^m_{it,S})^2 - (Y^A_{it} - \widehat{y}^m_{it,S+})^2, \qquad \bar{G}_{it} = \frac{1}{3}\sum_{m \in \{R,P,B\}} G^m_{it}. \tag{35}$$

The three procedures are ridge, persistence, and gradient boosting. A histogram boosting regression estimates $g^M(H_{it}) = E(\bar{G}_{it} \mid H_{it})$. Five candidate architectures vary the number of leaves, minimum leaf size, and L_2 penalty. The architecture for 2018 is selected by five-fold bank validation within target year 2017. For every later evaluation year, each candidate is fitted through two years before evaluation and scored on the immediately preceding target year. The selected candidate is then refitted through the preceding year. The resulting annual score uses prior-year outcomes and is applied before the intermediate report. The evaluation period is 2018–2025.

At each origin the acquisition rule selects the largest $\lfloor bn_t \rfloor$ values of $g^M(H_{it})$ for $b \in \{0.10, 0.25, 0.50\}$. It uses one score and the equal-weight loss across the three forecast procedures. Over 16,370 bank origins, the score reduces pooled loss by 12.98, 25.05, and 33.84 percent as the report budget rises. Uniform allocation reduces loss by 4.33, 10.86, and 21.73 percent. The mixed-value gains per eligible origin over uniform allocation are 3.962, 6.501, and 5.545 squared percentage points; simultaneous 95 percent lower endpoints across the three budgets are 1.521, 4.283, and 2.719. Deleting any one of the sixteen target quarters leaves every gain positive. Innovation-scale selection reduces pooled loss by 14.98, 22.26, and 31.44 percent. At the 25 percent budget, the mixed-value difference is 1.278 squared percentage points. Each simultaneous interval for the difference includes zero. Innovation scale therefore provides a strong information-only benchmark and accounts for much of the allocation value in the richer state. Figure 6 groups the annual ex ante score into within-year fifths. Realized paired gains are similar in the first two fifths and rise sharply across the upper three for every forecast procedure. The 10 percent budget

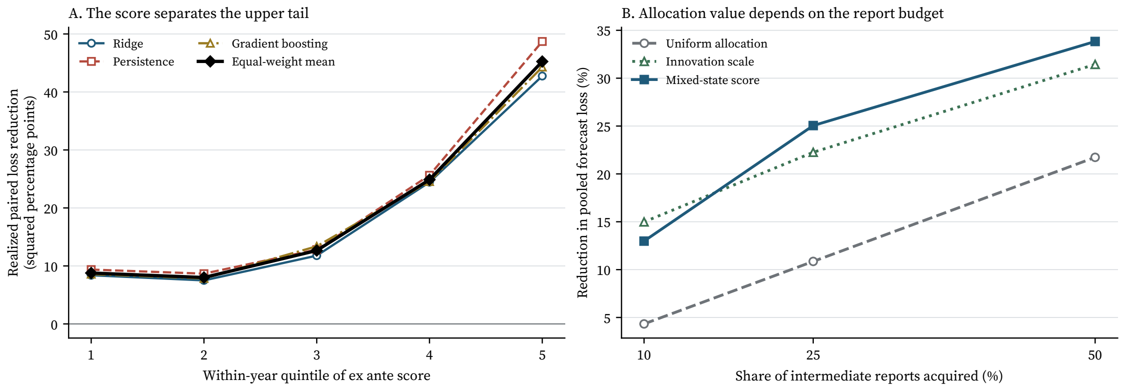

FIGURE 6. Score ordering and report-budget performance

The score estimates the equal-weight average paired gain across ridge, persistence, and gradient boosting using earlier target years. Panel A orders observations into within-year fifths and reports realized paired loss reduction for each procedure and their equal-weight mean. Panel B reports the percentage reduction in pooled forecast loss under uniform allocation, innovation-scale ranking, and the mixed-state score. The first two fifths in Panel A are close and locally reversed; the evidence for score ordering comes from the upper three fifths. At the 10 percent budget in Panel B, innovation scale yields the largest point estimate.

favors innovation scale, while the mixed-state score is higher at the 25 and 50 percent budgets; the corresponding differences from innovation scale remain unresolved at the reported precision. The acquisition exercise evaluates one downstream target under three mappings from records to predictions; variation in G_{it}^{m} supplies the component losses. With a monetary conversion ω_D dollars per squared percentage point of downstream loss and item cost c_{it}, the net acquisition value is $\omega_D g^M(H_{it}) - c_{it}$. An unconstrained planner acquires the report when this expression is nonnegative; a budgeted planner solves the associated binary knapsack problem. The Call Report data identify forecast-loss gains and report counts. The FFIEC burden figures in Table 2 measure whole-form hours. Item-level time, shared preparation costs, and ω_D determine a monetary acquisition policy.

The Gaussian archive score, assets, predictive variance, point-mass/change plug-in value, forecast-specific learners, and candidate tuning results are retained in the reproduction files. The public comparison reports the uniform benchmark, the strongest single-component ranking, and the sequential mixed-state score. This ordering asks whether the archive helps allocate reports at all, whether innovation scale explains the gain, and whether the additional state improves on that simpler rule. The intervals resolve the first comparison and leave the third open at the reported precision.

7.3. Changing the downstream target

Archive value depends on the target and on the records already retained. For next-quarter income, the additional ARM report contributes little once current earnings and maturity information are available; the intervals for changes in its marginal contribution include zero. A persistence forecast also outperforms the fitted ridge income forecast in 2016. Appendix K gives the estimates and ablations. These comparisons delimit the composition-target acquisition results and the prediction procedures used to value a report.

7.4. The Scope of the Empirical Correspondence

The reform sequence supplies observed changes in reporting frequency, form adoption paths, adjacent-quarter form persistence, estimated quarterly preparation hours, and a stated conversion-cost benchmark. These objects correspond to flow reporting costs, installed form choice, and reconfiguration resources. The threshold windows contain too few adopters for a reform-based causal estimate of adjustment curvature, so the burden accounting provides a scale while the curvature remains a model parameter.

The archive posterior, its propagation to future composition, and the expected value of an acquired report instantiate Lemma 4. The record-length and recency comparisons correspond to Proposition 2: the value of an old record depends on the state transition linking it to the target date. The mixed-state rule then maps the inherited history and current controls into a sequential estimate of the report's downstream loss reduction. Its budget frontier is evaluated against uniform and innovation-scale allocation.

The data identify report schedules, form choices, whole-form preparation hours, archive-conditioned forecast-loss gains, and the content of acquired reports. Mapping these quantities into the planner's monetary objective requires item-level costs, policy actions, and the conversion ω_D. Desired angular directions and response maps also remain model primitives. The reported accounting and break-even formula keep each unmeasured quantity visible in the welfare calculation.

8. Conclusion

A reporting system serves current operations and leaves an archive for later decisions. The model separates these uses. Specialization concentrates information along the operating priority, while transitional reports can preserve distinctions needed by later users. Inherited precision then changes the preferred direction of the next report. The planner's wedge and archive-value payment express this intertemporal information externality under the stated targets and resource costs.

The Call Report application measures a corresponding prediction margin. Retained histories determine how useful another report is for the declared income and composition targets. Prior-year selection of the sequential acquisition score gives positive gains over uniform allocation at each reported budget under simultaneous inference. Innovation scale captures much of the same value, and the richer state's incremental gain remains unresolved by those intervals. The comparison therefore supports the value of archive-conditioned acquisition while keeping the source of that value open.

Form-retention evidence, prediction gains, and welfare have separate inputs. Published whole-form hours describe a reporting-cost scale. The loss calculations use specified forecasts and eligible cohorts. Translating them into a monetary policy gain additionally requires item-level costs and the value attached to the target loss. The break-even calculation displays the item-level cost and target value needed to turn forecast accuracy into a monetary comparison.

Institutions consequently have a reporting choice that extends beyond the current task: which distinctions should the next archive preserve? The theory characterizes that choice within a specified information technology, and the empirical exercise shows how an existing archive can discipline a feasible acquisition comparison. Additional form-assignment evidence would connect the allocation mechanism to the effects of a reporting reform.

A. Target Information and Operating Directions

A.1. Response, Measurement, and Economic Units

Let $\theta \in \mathbb{R}^d$ collect economic coefficients with a fixed interpretation and unit for each coordinate. A local response map and an observation experiment have derivatives

$$G_r = V_r \Sigma W', \qquad K_r = U_r \Lambda U_r', \qquad \mathcal{I}_r = G_r' K_r G_r = W \Sigma R_r' \Lambda R_r \Sigma W', \quad R_r = U_r' V_r. \tag{36}$$

For the square specification, U_r, V_r, W are orthogonal, Σ contains response singular values, and Λ contains measurement-information eigenvalues. For example, $Y = H_r x + \varepsilon$, $x = G_r \theta$, and $\varepsilon \sim N(0, \Omega_r)$ give $K_r = H_r' \Omega_r^{-1} H_r$. The capacities Σ, Λ and the right basis W are held fixed across the comparisons that follow.

Fixed capacities give an isospectral structural experiment under the additional restriction $\Sigma = gI_d$, with $g > 0$. The analysis uses $W = I_d$ in the declared economic coordinates. A change of units transforms both the information matrix and the target gradient. The target cost below is invariant to that transformation. Spectral comparisons are made in the declared units, or relative to a declared positive-definite metric as described in Appendix A.

A.2. Fresh-Observation Information

For a differentiable target with nonzero gradient a, define

$$b_{\mathcal{I}}(a) = \begin{cases} a' \mathcal{I}^+ a, & a \in \operatorname{Range}(\mathcal{I}), \\ +\infty, & a \notin \operatorname{Range}(\mathcal{I}), \end{cases} \qquad \ell_{\mathcal{I}}(a) = b_{\mathcal{I}}(a)^{-1}. \tag{37}$$

The conventions are $1/(+\infty) = 0$ and $\mathcal{I}^+$ equal to the Moore–Penrose inverse.

LEMMA 5 (Target information). *For $\mathcal{I} \geq 0$ and $a \neq 0$, $\ell_{\mathcal{I}}(a) = \inf_{a'v=1} v' \mathcal{I} v$. This quantity and $b_{\mathcal{I}}(a)$ are invariant under an invertible change of parameter coordinates with the corresponding transformation of a.*

For a unit target $e(\varphi)$ and $\lambda_L > 0$,

$$b_\varphi(\delta) = \frac{1}{g^2} \left\{ \frac{\cos^2(\varphi - \delta)}{\lambda_H} + \frac{\sin^2(\varphi - \delta)}{\lambda_L} \right\}, \qquad \frac{1}{g^2 \lambda_H} \leq b_\varphi(\delta) \leq \frac{1}{g^2 \lambda_L}. \tag{38}$$

For d orthonormal targets, the set of such costs has a matrix characterization. Write $b \prec z$ when the sums of the largest k coordinates of b are bounded by those of z for $k < d$ and their total sums agree.

PROPOSITION 4 (Fresh-observation allocation). *Let* $\lambda_1 \ge \cdots \ge \lambda_d > 0$, $\mathcal{I}_R = g^2 R' \Lambda R$, *and let* A *be an orthogonal matrix whose columns are target gradients. As* R *ranges over the orthogonal group, the attainable vectors* $\operatorname{diag}(A'\mathcal{I}_R^{-1}A)$ *form*

$$\mathcal{B}(\Lambda) = \{b : b \prec g^{-2}(\lambda_d^{-1}, \ldots, \lambda_1^{-1})'\} = \operatorname{conv}\{\pi(g^{-2}\lambda^{-1}) : \pi \text{ is a permutation}\}. \tag{39}$$

For target-demand eigenvalues $d_1 \ge \cdots \ge d_d \ge 0$, *the minimum of* $\operatorname{tr}(D\mathcal{I}_R^{-1})$ *is* $g^{-2}\sum_j d_j/\lambda_j$.

The set in Proposition 4 is the Schur–Horn set of the inverse information matrix. Its trace is a fresh-observation accounting identity. Economic restrictions on the realized orientation enter through the decision problems in the next section. The pure target-demand minimization in the proposition sets operating and adjustment costs aside. A planner with positive operating costs retains those costs when its adjustment cost tends to zero.

A.3. Proof of Lemma 5

Suppose first that $a \in \operatorname{Range}(\mathcal{I})$. Decompose $v = v_R + v_N$ into the information range and kernel. The constraint is $a'v_R = 1$. On the range, Cauchy–Schwarz gives

$$1 = (a'v_R)^2 \le (a'\mathcal{I}^+ a)(v_R'\mathcal{I}v_R). \tag{40}$$

Equality is attained at $v_R = \mathcal{I}^+ a/(a'\mathcal{I}^+ a)$. Thus the infimum equals $(a'\mathcal{I}^+ a)^{-1}$. If a has a nonzero kernel projection a_N, the vector $v = a_N/\|a_N\|^2$ satisfies $a'v = 1$ and $v'\mathcal{I}v = 0$. This proves the range convention and the variational formula.

Let $\theta = T\vartheta$ with invertible T. The transformed information and target gradient are $\mathcal{I}_\vartheta = T'\mathcal{I}T$ and $a_\vartheta = T'a$. The substitution $v_\theta = Tv_\vartheta$ is a bijection between the two feasible tangent sets and preserves the quadratic form. The variational infima agree, so both target information and its inverse agree, including the zero and infinite cases.

A.4. Units and a Declared Metric

The raw eigenvalues of information depend on parameter units. For example, $\operatorname{diag}(4, 1)$ and $\operatorname{diag}(1, 4)$ have the same eigenvalues. Under $\theta = \operatorname{diag}(2, 1)\vartheta$, the transformed matrices are $\operatorname{diag}(16, 1)$ and $\operatorname{diag}(4, 4)$. Lemma 5 applies to the corresponding transformed gradients in each experiment.

A positive-definite tangent metric H gives an alternative way to fix the comparison. Use the generalized eigenvalues in $\mathcal{I}v = \lambda Hv$. Under the coordinate change, $H_\vartheta = T'HT$, and the generalized eigenvalues are preserved. Equivalently, use the eigenvalues of $H^{-1/2}\mathcal{I}H^{-1/2}$. Proposition 4 then applies in the whitened coordinates to a target basis satisfying $a_i'H^{-1}a_j =$

$\mathbf{1}\{i = j\}$, provided the maintained information family has the specified orthogonal orbit in those coordinates. The paper uses the declared economic units and $H = I$.

A.5. Operating Choices

Let $B = LSN'$ be a singular value decomposition. For an orthogonal V, $L'VN$ is orthogonal, and

$$\operatorname{tr}(B'V) = \operatorname{tr}(SL'VN) \leq \sum_j S_{jj}. \tag{41}$$

The value is attained at $V = LN'$. If all singular values are positive, equality requires every diagonal entry of $L'VN$ to equal one, which implies $L'VN = I$. Hence the maximizer is unique and equals $B(B'B)^{-1/2}$.

Write $C = L_C \operatorname{diag}(c_1, \dots, c_d) L_C'$ with $c_1 \geq \cdots \geq c_d$. For $Z = L_C'U$, the entries Z_{ij}^2 form a doubly stochastic matrix. The objective equals $\sum_{ij} c_i \lambda_j Z_{ij}^2$. The rearrangement inequality bounds it above by $\sum_j c_j \lambda_j$, attained by ordered eigenvector alignment. Repeated eigenvalues allow rotations within the corresponding eigenspaces. These static choices are conditional on the moments M_r, C_r; restrictions on their empirical variation must be supplied by the application.

A.6. Proof of Proposition 4

The inverse matrix is $\mathcal{I}_R^{-1} = g^{-2} R' \Lambda^{-1} R$. As R ranges over the orthogonal group, RA also ranges over that group. The real symmetric Schur–Horn theorem states that the diagonals of the orthogonal orbit of a matrix with ordered eigenvalues z are exactly the vectors majorized by z. This gives (39), with $z = g^{-2}(\lambda_d^{-1}, \dots, \lambda_1^{-1})'$. The convex-hull representation is the equivalent permutation characterization of majorization.

Diagonalize $D = L_D \operatorname{diag}(d_1, \dots, d_d) L_D'$. The trace $\operatorname{tr}(D\mathcal{I}_R^{-1})$ is minimized by pairing the largest d_j with the smallest inverse-information eigenvalue $1/(g^2\lambda_j)$. This yields the stated minimum. Pairing in the reverse order gives the maximum $g^{-2} \sum_j d_j/\lambda_{d+1-j}$.

In the rank-deficient case, the positive-information subspace rotates with the orthogonal basis. A target has finite fresh-observation cost precisely when its gradient belongs to that subspace. This condition uses all free components of the parameter. For example, consider

$$Y = \cos\delta\, \theta_1 + \sin\delta\, \theta_2 + \varepsilon, \qquad \varepsilon \sim N(0, 1). \tag{42}$$

For target θ_1 and $\sin\delta \neq 0$, $v = (1, -\cot\delta)'$ satisfies the target constraint and lies in the information kernel. Lemma 5 gives target information zero. If θ_2 is fixed and known, the scalar experiment has information $\cos^2\delta$ for θ_1. The scalar experiment in Appendix E specifies its parameter dimension.

B. The Institutional Control Problem

B.1. Proof of Lemma 1: Quadratic Value and the Priority Sequence

For $0 < \eta < 1$, write the date-t value for one institution as

$$V_t(x) = \frac{1}{2}Px^2 - r_t x + s_t. \tag{43}$$

The Bellman problem is

$$V_t(x) = \min_y \left\{ \frac{\rho}{2}(y - \bar{x}_t)^2 + \frac{\kappa}{2}(y - x)^2 + \eta V_{t+1}(y) \right\}. \tag{44}$$

Let $H_0 = \rho + \kappa + \eta P$. Completing the square yields

$$y_t(x) = \frac{\kappa x + \rho \bar{x}_t + \eta r_{t+1}}{H_0}, \qquad P = \kappa - \frac{\kappa^2}{H_0}, \qquad r_t = \frac{\kappa}{H_0}(\rho \bar{x}_t + \eta r_{t+1}). \tag{45}$$

The coefficient equation is the Riccati equation in (5). Its positive solution is the stated root of

$$\eta P^2 + [\rho + \kappa(1 - \eta)]P - \kappa\rho = 0. \tag{46}$$

It gives $0 < \gamma = \kappa/H_0 < 1$ and $P = \kappa(1 - \gamma)$. Bounded priorities and $\eta\gamma < 1$ give the bounded solution

$$r_t = \gamma\rho \sum_{j=0}^{\infty} (\eta\gamma)^j \bar{x}_{t+j}. \tag{47}$$

The constant term satisfies $s_t = \rho\bar{x}_t^2/2 + \eta s_{t+1} - (\rho\bar{x}_t + \eta r_{t+1})^2/(2H_0)$ and has a bounded solution obtained by discounted summation. Substituting (47) into the optimizer gives

$$y_t(x) = \gamma x + \frac{\rho}{H_0} \sum_{j=0}^{\infty} (\eta\gamma)^j \bar{x}_{t+j}. \tag{48}$$

The identity $\rho/H_0 = (1 - \gamma)(1 - \eta\gamma)$ transforms this expression into (6).

B.2. Optimality and Uniqueness

For a bounded installed path, iterating the completed-square Bellman identity and letting the terminal date grow gives its total cost as $V_0(x_0)$ plus

$$\frac{H_0}{2}\sum_{t=0}^{\infty}\eta^t[x_{t+1}-y_t(x_t)]^2. \tag{49}$$

The same identity applies to every finite-cost path: bounded priorities and the operating-cost term imply $\sum_t \eta^t x_{t+1}^2 < \infty$, hence the discounted quadratic terminal value tends to zero. The path in (6) is bounded and makes every square zero. The sum of squares is positive for any path differing at a finite date. This proves optimality and uniqueness for $0 < \eta < 1$. The per-date one-period problem gives $\gamma = \kappa/(\rho+\kappa)$ when continuation is assigned zero weight.

The two institutions share the same γ and averaging weights. Taking the difference of their policy functions gives the relative-direction law. Under a constant priority, iteration gives $\delta_t - \bar{\delta} = \gamma^t(\delta_0 - \bar{\delta})$. Finally,

$$\frac{\rho}{\kappa} = \frac{(1-\gamma)(1-\eta\gamma)}{\gamma} = \frac{1}{\gamma} - (1+\eta) + \eta\gamma. \tag{50}$$

The right side has derivative $-1/\gamma^2 + \eta < 0$. Increasing κ therefore increases γ. This establishes the comparative static in Lemma 1.

B.3. Unequal Coefficients and Matrix Installation

For institution $j \in \{A, O\}$, let $(\rho_j, \kappa_j, \eta_j)$ define its own Riccati solution and persistence γ_j. Its policy is (6) with those coefficients. The difference $\beta_{t+1} - \alpha_{t+1}$ depends on both inherited angles whenever the persistence coefficients differ. A state containing both installed angles then supplies the required inherited information.

For matrix-valued directions, installed bases can be updated by $V_{t+1} = V_t \exp(A_t)$ and $U_{t+1} = U_t \exp(B_t)$, where A_t, B_t are skew symmetric. Their relative orientation satisfies

$$R_{t+1} = \exp(-B_t) R_t \exp(A_t). \tag{51}$$

This identity records the effect of two installation choices. The solved control problem in Lemma 1 uses the scalar angular coordinate and the quadratic objective (4).

C. Posterior Loss and the Learning Limit

C.1. Proof of Lemma 2

Conditional Gaussian likelihoods add their quadratic terms to the log prior. Their Hessians sum to J_T in (7), giving posterior covariance J_T^{-1}. Completing the policy-loss square gives

$$L(q,\theta;a,c) - L(q^*(\theta),\theta;a,c) = \frac{c}{2}\left(q - \frac{a'\theta}{c}\right)^2. \tag{52}$$

Conditional minimization gives $\widehat{q} = a'\mu_T/c$. Its conditional expected loss is $a'J_T^{-1}a/(2c)$. Averaging over the independent policy problem gives (8). For fixed $\mathcal{I} > 0$,

$$(J_0 + n\mathcal{I})^{-1} = \frac{1}{n}\mathcal{I}^{-1} - \frac{1}{n^2}\mathcal{I}^{-1}J_0\mathcal{I}^{-1} + O(n^{-3}). \tag{53}$$

The trace expansion proves the stated approximation. The smallest eigenvalue of the fixed matrix $\mathcal{I}$ is positive throughout this calculation.

C.2. Proof of Theorem 1

In the orthonormal basis (u, v), put $z_s = \gamma^s\varepsilon_0$ and write

$$\begin{gathered}J_T = \begin{pmatrix} A_T & B_T \\ B_T & C_T \end{pmatrix}, \qquad A_T = u'J_0u + \tau\sum_{s<T}\cos^2 z_s, \\ B_T = u'J_0v + \tau\sum_{s<T}\cos z_s \sin z_s, \qquad C_T = v'J_0v + \tau\sum_{s<T}\sin^2 z_s.\end{gathered} \tag{54}$$

The bounds $|\sin z_s| \le |z_s|$ and $|\cos z_s| \le 1$ imply

$$\sum_{s=0}^{\infty}\sin^2 z_s \le \frac{\varepsilon_0^2}{1-\gamma^2}, \qquad \sum_{s=0}^{\infty}|\cos z_s \sin z_s| \le \frac{|\varepsilon_0|}{1-\gamma}. \tag{55}$$

Thus $C_T \to C_\infty > 0$, B_T has a finite limit, and $A_T = \tau T + O(1)$. Let $S_T = C_T - B_T^2/A_T$. Then $S_T \to C_\infty$, and block inversion gives

$$J_T^{-1} = \begin{pmatrix} A_T^{-1} + B_T^2/(A_T^2 S_T) & -B_T/(A_T S_T) \\ -B_T/(A_T S_T) & S_T^{-1} \end{pmatrix} \longrightarrow \begin{pmatrix} 0 & 0 \\ 0 & C_\infty^{-1} \end{pmatrix}. \tag{56}$$

Returning to the original coordinates gives (10). Taking the trace against $D/2$ gives the limiting policy loss. If $0 < |\varepsilon_0| \le \pi/2$, each term with $s \ge 1$ strictly increases with γ because $\sin^2 z$ increases with $|z|$ for $|z| \in (0, \pi/2)$. The sum is therefore strictly increasing. Combining this with (50) proves the statement about adjustment costs.

C.3. Proof of Proposition 1

Substituting $\delta_t - \bar{\delta} = \gamma^t \varepsilon_0$ into the relative operating and adjustment costs and summing their discounted values gives (11). Its derivative has the sign of $\rho\gamma - \kappa(1-\gamma)(1-\eta\gamma)$. This expression is strictly increasing from $-\kappa$ to ρ, so it has one root $\gamma^E \in (0,1)$, and $\mathcal{C}^G$ decreases below that root and increases above it. Its second derivative is positive at γ^E.

Theorem 1 makes $C_\infty(\gamma)$ strictly increasing on $(0,1)$, so $\mathcal{L}^G$ is strictly decreasing there. As $\gamma \uparrow 1$, the first M terms in (9) converge to $M\sin^2(\varepsilon_0)$ for each fixed M; letting M grow gives $C_\infty(\gamma) \to \infty$. The extensions of $\mathcal{C}^G$ and $\mathcal{L}^G$ are therefore continuous at one. On $[0,\gamma^E]$, the social objective decreases strictly, including at γ^E through the policy-loss derivative. Compactness then gives a global minimizer in $(\gamma^E, 1]$.

At $\chi = 0$, the unique global minimizer is the interior point γ^E. Continuity of the objective confines every minimizer to a neighborhood of γ^E for sufficiently small χ. The implicit-function theorem applied to $\mathcal{C}^{G\prime}(\gamma) + \chi\mathcal{L}^{G\prime}(\gamma) = 0$ supplies a unique local minimizer in that neighborhood. Since $\mathcal{L}^{G\prime}(\gamma^E) = -\bar{d}C_\infty'(\gamma^E)/[2C_\infty(\gamma^E)^2]$, differentiation at $\chi = 0$ gives (13). Positivity follows from $C_\infty'(\gamma^E) > 0$ and $\mathcal{C}^{G\prime\prime}(\gamma^E) > 0$.

C.4. Proof of Corollary 1

With $\lambda_L > 0$, each batch satisfies $n\mathcal{I}(\delta_s) \succeq ng^2\lambda_L I_2$. Therefore $J_T \succeq Tng^2\lambda_L I_2$ and $J_T^{-1} \preceq (Tng^2\lambda_L)^{-1} I_2$. Taking the trace against $D/2$ proves the bound.

In the rank-one case, the supplemental observations add ξm_T to C_T and leave A_T, B_T unchanged in (54). The Schur complement is $S_T = C_T + \xi m_T - B_T^2/A_T$. Along the path of Theorem 1, $A_T \to \infty$, B_T stays bounded, and C_T converges to C_∞. Hence $S_T \to \infty$ if and only if $m_T \to \infty$. Under this condition, every entry of the inverse in (56) tends to zero. Otherwise the nondecreasing cumulative counts are bounded, and the vv' posterior variance stays bounded away from zero. This proves necessity and sufficiency. The argument applies to every diverging sequence m_T, including sequences with $m_T/T \to 0$.

C.5. Precision, Rank, and Policy Targets

For positive semidefinite information matrices I_s,

$$\ker\left(\sum_s I_s\right) = \bigcap_s \ker(I_s). \tag{57}$$

Indeed, $v'(\sum_s I_s)v = \sum_s v' I_s v$ vanishes precisely when every nonnegative summand vanishes. This proves the finite-archive rank statement in the main text. Posterior concentration in every direction requires the smallest eigenvalue of J_T to diverge. Theorem 1 computes that eigenvalue's finite limit under the specified path.

For a restricted set of future targets, the posterior loss is $\mathrm{tr}(DJ_T^{-1})/2$. In Theorem 1, its limit is zero exactly when $v'Dv = 0$. Because $D \geq 0$, this condition is equivalent to future target gradients lying in u's span with probability one under their loss weights. The policy consequence of residual parameter uncertainty thus depends on the target distribution.

D. The Planner and Reporting Threshold

D.1. State Reduction and Existence

Let $\alpha = m - \delta/2$ and $\beta = m + \delta/2$, and define the corresponding common and relative desired positions. The sum of operating costs is $\rho(m_{t+1} - \bar{m}_t)^2 + \rho(\delta_{t+1} - \bar{\delta}_t)^2/4$. The sum of adjustment costs is $\kappa(m_{t+1} - m_t)^2 + \kappa(\delta_{t+1} - \delta_t)^2/4$. Information in (3) depends on the relative angle. Minimizing over the common path therefore adds a term independent of all relative choices and yields (20).

For finite H, the value recursion starts with the continuous zero function. At any positive-definite J, policy risk is finite, continuous, and nonnegative. By induction, the continuation value is continuous and nonnegative. On compact sets of states (x, J) with $J > 0$, the quadratic term $\kappa(y - x)^2/4$ supplies a uniform bound on minimizing choices, using the feasible choice $y = x$ as a locally bounded upper bound for value. Minimization over that compact set gives a continuous value and an attained minimum. This establishes existence by backward induction.

For $J_2 \geq J_1 > 0$, any fixed future angle sequence has weakly smaller policy losses starting from J_2. The sequence's operating costs are unchanged. Taking the infimum over sequences gives $\mathcal{V}_t(x, J_2) \leq \mathcal{V}_t(x, J_1)$. Where the matrix gradient exists under the trace inner product, it is negative semidefinite. Hence $M_t \geq 0$ in (22).

D.2. The Marginal Reporting Condition

At a differentiability point of the continuation value, perturb y and write $J^+ = J + n_t\mathcal{I}(y)$. The derivative of inverse precision is $-n_t(J^+)^{-1}\mathcal{I}'(y)(J^+)^{-1}$. The chain rule applied to (21) yields

$$\begin{aligned} 0 = {} & \frac{\rho}{2}(y - \bar{\delta}_t) + \frac{\kappa}{2}(y - x) - \frac{\chi n_t}{2}\,\mathrm{tr}\{(J^+)^{-1}D_t(J^+)^{-1}\mathcal{I}'(y)\} \\ & + \eta\partial_x\mathcal{V}_{t+1}(y, J^+) + \eta n_t\,\mathrm{tr}\{\nabla_J\mathcal{V}_{t+1}(y, J^+)\mathcal{I}'(y)\}. \end{aligned} \tag{58}$$

Collecting the two information derivatives gives (23). At a state with multiple optimal continuation branches, the Bellman minimization remains the definition of value; the displayed derivative condition applies to smooth branches where the indicated derivatives exist.

D.3. Proof of Theorem 2

Let $A = \mathrm{diag}(a_0, b_0)$. The posterior precision is $A + \tau e(\delta)e(\delta)'$. Its trace is N and its determinant is

$$\det[A + \tau e(\delta)e(\delta)'] = a_0 b_0 + \tau\{b_0 \cos^2\delta + a_0 \sin^2\delta\} = K + L\sin^2\delta. \tag{59}$$

For a two-by-two positive-definite matrix, the inverse trace is its trace divided by its determinant. This gives (29). If $a_0 = b_0$, then $L = 0$ and the information term is constant, so the quadratic cost selects zero.

Suppose $a_0 > b_0$. Then $L > 0$. From (32),

$$F(\delta) - F(0) \geq \left(\frac{h}{4} - \frac{\chi NL}{2K^2}\right)\delta^2. \tag{60}$$

The bound follows by $K + L\sin^2\delta \geq K$ and $\sin^2\delta \leq \delta^2$. It is strict for a nonzero δ at $\chi = \chi^*$ because $\sin^2\delta < \delta^2$. Thus zero is the unique global minimum up to and including the threshold.

For $\chi > \chi^*$, $F''(0) = h/2 - \chi NL/K^2 < 0$. For $0 < \delta < \pi/2$, a stationary point solves $\chi = \Xi(\delta)$, where

$$\Xi(\delta) = \frac{h\delta[K + L\sin^2\delta]^2}{NL\sin(2\delta)}. \tag{61}$$

The function $\delta/\sin(2\delta)$ strictly increases on this interval: its derivative has the sign of $\sin(2\delta) - 2\delta\cos(2\delta)$, whose derivative is $4\delta\sin(2\delta) > 0$ and whose limit at zero is zero. The factor $[K + L\sin^2\delta]^2$ is positive and strictly increasing. Therefore Ξ strictly increases, with limits χ^* at zero and $+\infty$ at $\pi/2$. There is exactly one positive stationary point, the objective decreases before it and increases after it, and reflection gives the negative minimizer. Periodicity of the information term and the displacement cost exclude every representative with larger absolute angle. This proves the global characterization and the limits of δ^*.

To differentiate the threshold, set $z = a_0 - b_0 > 0$, $d = b_0 + \tau$, and $e = 2b_0 + \tau$. Then

$$\chi^* = \frac{hb_0^2}{2\tau}\frac{(z+d)^2}{z(z+e)}, \qquad \frac{d}{dz}\log\frac{(z+d)^2}{z(z+e)} = \frac{-\tau z - de}{(z+d)z(z+e)} < 0. \tag{62}$$

This proves the derivative and both limits in the theorem. An archive $J = j_2 I_2 + p e_1 e_1'$ with $p \geq 0$ provides the comparison directly: increasing the precision p of retained first-coordinate reports raises a_0 while leaving b_0 fixed.

D.4. Proof of Corollary 2

With a constant zero priority, the institution's value at an installed angle x is $Px^2/2$. After the planner chooses (α', β'), discounted operating continuation is $\eta P(\alpha'^2 + \beta'^2)/2$. Write $m' = (\alpha' + \beta')/2$ and $\delta = \beta' - \alpha'$. Starting from zero and adding the current operating and adjustment costs gives $h_c m'^2 + h_c \delta^2/4$, where $h_c = \rho + \kappa + \eta P$. The common direction minimizes at zero. The policy loss of the next report is unchanged, so the objective is (29) with curvature h_c. Theorem 2 applies with that substitution.

D.5. The Terminal Local Change

Let $F_\chi(\delta) = h(\delta - \delta^E)^2/4 + \chi r(\delta)$. Since inherited $J > 0$, r is smooth. The first-order condition is $h(\delta - \delta^E)/2 + \chi r'(\delta) = 0$. At $(\delta, \chi) = (\delta^E, 0)$ its derivative with respect to δ is $h/2 > 0$. The implicit function theorem gives $\delta^{SP} - \delta^E = -2\chi r'(\delta^E)/h + O(\chi^2)$. Substituting this into a third-order expansion of F_χ gives the objective gain in (33). The positive cost curvature makes the stationary branch a local minimum for sufficiently small χ.

E. A Scalar Experiment at an Information Boundary

This appendix uses one free structural parameter $\theta \in \mathbb{R}$ and a two-dimensional response. Let $x = ge(\alpha)\theta$ and observe

$$Y = \sqrt{\lambda}\, e(\beta)' x + \varepsilon = g\sqrt{\lambda}\cos(\beta - \alpha)\theta + \varepsilon, \qquad \varepsilon \sim N(0, 1). \tag{63}$$

The response derivative is $G = ge(\alpha) \in \mathbb{R}^{2\times 1}$. Measurement information is $K = \lambda e(\beta)e(\beta)'$. The scalar information is $\mathcal{I}(r) = \lambda g^2 \cos^2 \delta(r)$. The response singular value and the positive measurement eigenvalue are fixed. The scalar information varies with their relative direction.

Suppose an announced regime r is expected to persist. The inherited direction is held fixed, and the preferred relative direction is $\bar{\delta}(r)$. Lemma 1 gives the next direction as $\delta(r) = \gamma\delta_{\text{old}} + (1-\gamma)\bar{\delta}(r)$. Let $\delta(r^*) \in \pi/2 + \pi\mathbb{Z}$, and suppose the first nonzero derivative of $\bar{\delta}$ at r^* has order $m \geq 1$. Taylor expansion gives

$$\mathcal{I}(r) = C_m|r - r^*|^{2m} + o(|r - r^*|^{2m}), \qquad C_m = \lambda g^2(1-\gamma)^2\left[\frac{\bar{\delta}^{(m)}(r^*)}{m!}\right]^2. \tag{64}$$

For $r_n - r^* = zn^{-1/(2m)}$, the total information from n fresh observations tends to $C_m z^{2m}$. This is a rate at which the regime approaches the boundary. With a scalar Gaussian prior of precision $j_0 > 0$, posterior variance along this sequence tends to $(j_0 + C_m z^{2m})^{-1}$. The order m is an assumption on the regime-priority schedule in this experiment.

The expansion follows from $\delta(r) - \delta(r^*) = (1-\gamma)\bar{\delta}^{(m)}(r^*)(r-r^*)^m/m! + o(|r-r^*|^m)$ and $\cos^2(\pi/2+z) = z^2 + o(z^2)$. It concerns the scalar parameter specified in (63). The rank-one two-parameter target in (42) uses the range condition in Lemma 5. In a nonlinear experiment, zero local Fisher information can also coexist with global injectivity; $Y \sim N(\theta^3, 1)$ at $\theta = 0$ is one example.

E.1. Projection Inference

Let an empirical moment vector $\widehat{m}_n$ have mean $m(\theta, r) \in \mathbb{R}^q$. Assume a uniform central limit theorem over the maintained parameter-regime class and uniform consistency of a covariance estimator $\widehat{\Omega}_n$, with covariance eigenvalues bounded above and bounded away from zero. Define

$$Q_n(\theta, r) = n[\widehat{m}_n - m(\theta, r)]'\widehat{\Omega}_n^{-1}[\widehat{m}_n - m(\theta, r)], \qquad \mathcal{C}_n(r) = \{\theta : Q_n(\theta, r) \le \chi^2_{q,1-\alpha}\}. \tag{65}$$

At the true parameter, the uniform central limit theorem and covariance consistency imply a uniform χ^2_q limit for Q_n. Hence $\liminf_n \inf_{(\theta,r)} P_{\theta,r}\{\theta \in \mathcal{C}_n(r)\} \ge 1-\alpha$. Projection gives the same lower coverage bound for $\{\psi(\vartheta) : \vartheta \in \mathcal{C}_n(r)\}$. For Gaussian sample means with known covariance, the true-parameter statistic has its chi-square distribution at each finite n. These assumptions concern moment distributions and covariance estimation; they allow the derivative of m with respect to θ to approach zero along the scalar regime sequence.

F. Record Labels and Observation Maps

F.1. Retaining the Observation Map

Let L be a recorded regime label with a probability distribution independent of θ. Let $s_L(Y;\theta)$ be the conditional score and assume the differentiability and integrability needed to interchange differentiation and integration. Integrating over the label gives score $E[s_L(Y;\theta) \mid Y]$. The matrix law of total variance gives

$$\mathcal{I}_{Y,L} - \mathcal{I}_Y = E[\mathrm{Var}\{s_L(Y;\theta) \mid Y\}] \ge 0. \tag{66}$$

The labelled archive in (7) keeps the observation maps associated with each report. An archive with missing maps requires the corresponding observed-data likelihood. Equation (66) describes the information difference for the stated mixture experiment.

G. Cumulative Excitation

G.1. Proof of Lemma 3

Let $S_T = \sum_{s<T} n_s \mathcal{I}(\delta_s) \succeq 0$. For a positive-definite matrix, $\|J_T^{-1}\|_{\mathrm{op}} = 1/\lambda_{\min}(J_T)$. Weyl's inequalities give $\lambda_{\min}(S_T) + \lambda_{\min}(J_0) \le \lambda_{\min}(J_T) \le \lambda_{\min}(S_T) + \lambda_{\max}(J_0)$. Thus the smallest eigenvalue of J_T diverges exactly when that of S_T does. Disjoint blocks of length L under the displayed finite-window condition give $S_T \succeq \lfloor T/L \rfloor \epsilon I_d$, proving the rate in Lemma 3. This is a covariance-concentration statement for the specified Gaussian posterior.

References

Sahil Ahmad, He Helen Huang, and Angela J. Yu. Cost-sensitive Bayesian control policy in human active sensing. *Frontiers in Human Neuroscience*, 8:955, 2014. doi: 10.3389/fnhum.2014.00955.

Ingvild Almås, Orazio Attanasio, and Pamela Jervis. Presidential address: Economics and measurement: New measures to model decision making. *Econometrica*, 92(4):947–978, 2024. doi: 10.3982/ECTA21528.

Donald W. K. Andrews and Patrik Guggenberger. Identification- and singularity-robust inference for moment condition models. *Quantitative Economics*, 10(4):1703–1746, 2019. doi: 10.3982/QE1219.

Isaiah Andrews and Anna Mikusheva. Conditional inference with a functional nuisance parameter. *Econometrica*, 84(4):1571–1612, 2016. doi: 10.3982/ECTA12868.

E. W. Bai and S. S. Sastry. Persistency of excitation, sufficient richness and parameter convergence in discrete time adaptive control. *Systems & Control Letters*, 6(3):153–163, 1985. doi: 10.1016/0167-6911(85)90035-0.

John S. Baras and Alain Bensoussan. Optimal sensor scheduling in nonlinear filtering of diffusion processes. *SIAM Journal on Control and Optimization*, 27(4):786–813, 1989. doi: 10.1137/0327042.

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.

Stephen Boyd and Shankar Sastry. On parameter convergence in adaptive control. *Systems & Control Letters*, 3(6):311–319, 1983. doi: 10.1016/0167-6911(83)90071-3.

Kathryn Chaloner and Isabella Verdinelli. Bayesian experimental design: A review. *Statistical Science*, 10(3):273–304, 1995. doi: 10.1214/ss/1177009939.

Gilles Chemla and Christopher Hennessy. Equilibrium counterfactuals. *International Economic Review*, 62(2):639–669, 2021. doi: 10.1111/iere.12513.

Herman Chernoff. Sequential design of experiments. *Annals of Mathematical Statistics*, 30(3): 755–770, 1959.

Mathias Dewatripont, Ian Jewitt, and Jean Tirole. The economics of career concerns, part II: Application to missions and accountability of government agencies. *Review of Economic Studies*, 66(1):199–217, 1999. doi: 10.1111/1467-937X.00085.

Mehmet Ekmekci, Leandro Gorno, Lucas Maestri, Jian Sun, and Dong Wei. Learning from manipulable signals. *American Economic Review*, 112(12):3995–4040, 2022. doi: 10.1257/aer.20211158.

Juan Carlos Escanciano. Semiparametric identification and fisher information. *Econometric Theory*, 38(2):301–338, 2022. doi: 10.1017/S0266466621000116.

Federal Deposit Insurance Corporation. BankFind suite: Financial, institution, and history data, 2026. URL `https://api.fdic.gov/banks/docs/`. API data and field definitions, downloaded August 28.

Federal Financial Institutions Examination Council. FFIEC 051: Data elements with reporting changes, 2016. URL
`https://www.ffiec.gov/sites/default/files/data/reporting-`

forms/FFIEC051_20161230_DataElementsImpact.pdf. December 30; ARM reporting frequency on p. 13.

Federal Financial Institutions Examination Council. FFIEC finalizes June 2017 proposed revisions to streamline the Call Report, 2018. URL https://www.ffiec.gov/news/press-releases/2018/pr-01-03. January 3.

Federal Financial Institutions Examination Council. Annual report 2025, 2026. URL https://www.ffiec.gov/sites/default/files/data/reports/annual-reports/FFIEC_AnnualReport_2025_WebLocked.pdf. Call Report full-review discussion.

Federal Reserve Bank of St. Louis. FRED: DFF, DGS2, and DGS10, 2026. URL https://fred.stlouisfed.org/series/DGS2. Daily federal funds and Treasury constant-maturity rates, downloaded August 28.

Janis Hardwick and Quentin F. Stout. Response adaptive designs that incorporate switching costs and constraints. *Journal of Statistical Planning and Inference*, 137(8):2654–2665, 2007. doi: 10.1016/j.jspi.2006.04.013.

Bengt Holmstrom and Paul Milgrom. Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design. *Journal of Law, Economics, & Organization*, 7(Special Issue): 24–52, 1991.

Roger A. Horn and Charles R. Johnson. *Matrix Analysis*. Cambridge University Press, Cambridge, 2 edition, 2012. doi: 10.1017/CBO9781139020411.

Xun Huan, Jayanth Jagalur, and Youssef Marzouk. Optimal experimental design: Formulations and computations. *Acta Numerica*, 33:715–840, 2024. doi: 10.1017/S0962492924000023.

Philippe Jehiel and Jakub Steiner. Selective sampling with information-storage constraints. *Economic Journal*, 130(630):1753–1781, 2020. doi: 10.1093/ej/uez068.

Tetsuya Kaji. Theory of weak identification in semiparametric models. *Econometrica*, 89(2): 733–763, 2021.

Rudolf E. Kalman. A new approach to linear filtering and prediction problems. *Journal of Basic Engineering*, 82(1):35–45, 1960. doi: 10.1115/1.3662552.

J. Kiefer and J. Sacks. Optimal sequential inference and design. *Annals of Mathematical Statistics*, 34(3):705–750, 1963.

Michael Kolonko and Heinrich Benzing. The sequential design of Bernoulli experiments including switching costs. *Operations Research*, 33(2):412–426, 1985. doi: 10.1287/opre.33.2.412.

H. W. J. Lee, K. L. Teo, and Andrew E. B. Lim. Sensor scheduling in continuous time. *Automatica*, 37(12):2017–2023, 2001. doi: 10.1016/S0005-1098(01)00159-5.

Alex S. Leong, Subhrakanti Dey, and Daniel E. Quevedo. Sensor scheduling in variance based event triggered estimation with packet drops, 2016. URL https://arxiv.org/abs/1511.04792. arXiv:1511.04792, version 4.

Wei Li. Changing one's mind when the facts change: Incentives of experts and the design of reporting protocols. *Review of Economic Studies*, 74(4):1175–1194, 2007. doi: 10.1111/j.1467-937X.2007.00446.x.

Annie Liang and Erik Madsen. Data and incentives. *Theoretical Economics*, 19(1):407–448, 2024. doi: 10.3982/TE5289.

Annie Liang, Xiaosheng Mu, and Vasilis Syrgkanis. Dynamically aggregating diverse information. *Econometrica*, 90(1):47–80, 2022. doi: 10.3982/ECTA18324.

Robert E. Jr. Lucas. Econometric policy evaluation: A critique. *Carnegie-Rochester Conference Series on Public Policy*, 1:19–46, 1976. doi: 10.1016/S0167-2231(76)80003-6.

Albert W. Marshall, Ingram Olkin, and Barry C. Arnold. *Inequalities: Theory of Majorization and Its Applications*. Springer, New York, 2 edition, 2011. doi: 10.1007/978-0-387-68276-1.

Michael McLeay and Silvana Tenreyro. Optimal inflation and the identification of the phillips curve. *NBER Macroeconomics Annual*, 34(1):199–255, 2020. doi: 10.1086/707181.

Manuel Oechslin and Elias Steiner. Statistical capacity and corrupt bureaucracies. *Review of International Organizations*, 17(1):143–174, 2022. doi: 10.1007/s11558-021-09421-5.

Office of the Comptroller of the Currency, Federal Reserve Board, and Federal Deposit Insurance Corporation. Agency information collection activities: Submission for OMB review; joint comment request, 2017. URL `https://www.govinfo.gov/content/pkg/FR-2017-01-09/pdf/2017-00085.pdf`. Federal Register 82, January 9, pp. 2444–2467.

Friedrich Pukelsheim. *Optimal Design of Experiments*. Society for Industrial and Applied Mathematics, Philadelphia, 2006. doi: 10.1137/1.9780898719109.

Maytal Saar-Tsechansky, Prem Melville, and Foster Provost. Active feature-value acquisition. *Management Science*, 55(4):664–684, 2009. doi: 10.1287/mnsc.1080.0952.

Laura Veldkamp and Cindy Chung. Data and the aggregate economy. *Journal of Economic Literature*, 62(2):458–484, 2024. doi: 10.1257/jel.20221580.

P. Whittle. Some general results in sequential design. *Journal of the Royal Statistical Society: Series B (Methodological)*, 27(3):371–387, 1965. doi: 10.1111/j.2517-6161.1965.tb00600.x.

Weijie Zhong. Optimal dynamic information acquisition. *Econometrica*, 90(4):1537–1582, 2022. doi: 10.3982/ECTA17787.

Ingvar Ziemann and Henrik Sandberg. On uninformative optimal policies in adaptive LQR with unknown B-matrix. In *Proceedings of the 3rd Conference on Learning for Dynamics and Control*, volume 144 of *Proceedings of Machine Learning Research*, pages 213–226. PMLR, 2021.